

about the world, immediately and without much conscious reflection. Our linguistic capacity thereby enables us to engage in complex cooperative activities that require rapid information

between large numbers of people, such as building bridges (or supercooled lakes that freeze but not unfortunately also waging war). We have our ability to perceive visually with many other species. In contrast, no other species has our advanced linguistic capacity.

Investigation into what makes humans distinctive as a species ends with language-use. The inquiry into the nature of linguistic capacities is therefore at the very center of the human sciences. Perhaps the most important progress in understanding our capacity for gaining and conveying information about the world linguistically has been slower progress in understanding (say) our perceptual capacities. We have made many inroads in grammar and phonology. But a complete account of the relation between an utterance and the information about the world it conveys is still probably no closer than a complete account of the biology of the species. A complex phenomenon needs to be broken up into many component parts to have any hope of ultimately being explained. The phenomenon of language-use is no different in this regard. My proposal is a contribution to only a small part of an account of the relation between utterances and the information about the world they convey. For example, there are very difficult questions about

# Language

in

# Context

# Selected Essays

Jason

Stanley

# Language in Context

*This page intentionally left blank*

# Language in Context

*Selected Essays*

Jason Stanley

CLARENDON PRESS · OXFORD

# OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford ox2 6dp

Oxford University Press is a department of the University of Oxford.

It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi

Kuala Lumpur Madrid Melbourne Mexico City Nairobi

New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece

Guatemala Hungary Italy Japan Poland Portugal Singapore

South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press  
in the UK and in certain other countries

Published in the United States

by Oxford University Press Inc., New York

© Jason Stanley 2007

The moral rights of the authors have been asserted

Database right Oxford University Press (maker)

First published 2007

All rights reserved. No part of this publication may be reproduced,  
stored in a retrieval system, or transmitted, in any form or by any means,  
without the prior permission in writing of Oxford University Press,  
or as expressly permitted by law, or under terms agreed with the appropriate  
reprographics rights organization. Enquiries concerning reproduction  
outside the scope of the above should be sent to the Rights Department,  
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover  
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

Typeset by Laserwords Private Limited, Chennai, India

Printed in Great Britain

on acid-free paper by

Biddles Ltd., King's Lynn, Norfolk

ISBN 9780199225927

ISBN 9780199225934 (Pbk.)

10 9 8 7 6 5 4 3 2 1

## PREFACE

My editor, Peter Momtchiloff, tells me that in the spring of 1995 I approached him about writing a book on the subject of context-dependence. So why a collection of papers rather than a book? At the time at which I planned such a book, I had just emerged from graduate school. Book-length treatments of topics are unwise for unknown junior professors. The strategy I decided to follow was to write a series of papers as if they were chapters of a book. If the papers attracted interest, I would use them as a basis for a book.

The initial plan was sensible. But philosophy as I practice it is a collective enterprise. I am deeply affected by interchanges with my professional colleagues. As a result, some of the papers ended up as joint ventures. The first argument against publishing this work as a book is that to do so would be to diminish the role my co-authors, Zoltán Gendler Szabó and then Jeffrey King, played in the construction of my overall view. The second argument for publishing this as a series of papers is that the individual papers were initially designed as chapters of a book. Even as late as 2004, when I wrote “Semantics in Context”, intending it as the penultimate chapter, I intended to rewrite the papers into coherent chapters. So the papers already possess the coherence of a book. They are collectively different sections in the construction of a specific metaphysical account of how context interacts with word meaning and sentence structure to create content.

Since I started working on this topic as a graduate student, the topic of the relation between context and linguistic content has moved to center stage in philosophy. A number of philosophers and cognitive scientists, dissatisfied with standard semantics and Gricean pragmatics, have proposed significant re-evaluations of the conceptual terrain. Within philosophy of language, a particularly popular strategy has been to diminish the role of semantics in an account of communication. There are various sociological explanations for the popularity of this view, such as the fact that well-known semantic theories of proper names required there to be a large gap between the semantic content of certain sentences and our ordinary intuitions about what is said. Unfortunately, philosophers who advocate a reduced role for semantic explanations have rarely provided explicit pragmatic theories that can replace their empirical coverage. The results have too often been views that advance negative conclusions about the scope of the theory of meaning, without supplementing the role of the theory of meaning with accounts that yield testable predictions. In fact, every single

book written by a philosopher or a cognitive scientist on the topic of context and communication in recent years has been principally devoted to negative assessments of semantic explanations, coupled with hopelessly underdeveloped sketches of pragmatic accounts. If one were to judge the topic of context and communication by the books on the topic, one would therefore think that the only dispute is between which underdeveloped pragmatic account of what is said is to be preferred. My central motivation in collecting my papers on this topic into a book is to provide a counterweight to this growing literature, and restore balance to the debate.

In the summer of 1997 Zoltán Gendler Szabó and I taught a seminar together on context-dependence at the Cornell Summer Linguistics Institute. The didactic nature of our joint paper “On Quantifier Domain Restriction” is explained by the fact that it was originally prepared as the template for this seminar. At the institute, Angelika Kratzer was giving a series of lectures drawing on classic work by Murvet Enc and Barbara Partee in which tenses and modals were treated in terms of explicit object-language quantification over times and worlds. I was greatly affected by the underlying picture of the syntax-semantics interface she was sketching. Kratzer’s view influenced the initial drafts of “Context and Logical Form”, which were written during this time and presented at my seminar with Zoltán. During my time at Cornell I also taught two seminars on context-dependence, in the fall of 1996 and the fall of 1999. Jeffrey King was visiting Cornell in the fall of 1999, and both he and Zoltán attended my seminar. I presented a great deal of the material in my subsequently published papers at this seminar, and their input had a great effect on the development of my views. The contributions of various participants of the seminar were also very helpful, most memorably Sally McConnell-Ginet, Susanna Siegel and Mandy Simons. During my time at Cornell there were also several conferences on context-dependence that were germane to the development of my work. For example, in 1998, Sally McConnell-Ginet, Zoltán Gendler Szabó, and I co-organized a conference on context-dependence, which brought together a number of researchers. This conference also had an influence on my subsequent work.

In the creation of the positive part of this work, I was influenced by the work of linguists and philosophers who had developed views on the semantic content of particular context-sensitive constructions. The authors of work which most affected my views (aside from my co-authors) are Robin Cooper, Max Cresswell, Kai von Stechow, Angelika Kratzer, Peter Ludlow, and Dag Westerstahl. After the first few papers were published, I began to engage with a different group, namely advocates of “truth-conditional pragmatics” such as François Recanati, as well as advocates of relevance theoretic pragmatics, such as Deirdre Wilson, Dan Sperber, and Robyn Carston. I first met this inter-disciplinary group when they invited me to a cognitive science conference in Oxford in the fall of 2000. At that conference, I presented the first draft of “Making it Articulated”, which was substantially improved by the discussion

at the event. I again had a chance to return to my concerns with relevance-theoretic explanations of apparently semantic phenomena in an interchange with Robyn Carston at the American Philosophical Association meetings in December of that year. For the next three years, I spent many days in conferences with members of the inter-disciplinary “semantics–pragmatics” group. Several papers have resulted from my attempts to respond to the powerful objections raised by advocates of non-semantic accounts of the data I discuss. Richard Breheny helped me both with his enjoyably aggressive verbal input, and his well-developed series of objections to my views that I respond to in “Semantics in Context”. Of all of the advocates of pragmatic approaches, I need to single out Kent Bach, Herman Cappelen, Robyn Carston, Ernest Lepore, and François Recanati for special mention. Though we have all annoyed each other tremendously with our obstinacies over the past five to ten years, I think it has all been (almost) worth it.

In addition to those mentioned above, a number of linguists and philosophers have been fruitful sources of insight in casual discussion and spirited debate. Many of these scholars are part of the semantics group brought together annually since 1997 by Ernie Lepore at Rutgers University, which is now (fittingly) my home institution. I would particularly like to thank Emma Borg, Keith DeRose, Michael Fara, Carl Ginet, Michael Glanzberg, John Hawthorne, Richard Heck, James Higginbotham, Pauline Jacobson, Christopher Kennedy, Hanna Kim, Peter Ludlow, Stephen Neale, Michael Nelson, Jeff Pelletier, Mark Richard, Roger Schwarzschild, Adam Sennet, Robert Stalnaker, Joseph Stern, Matthew Stone, Anna Szabolcsi, Richmond Thomason, and Timothy Williamson. My occasional discussions of philosophy with Delia Graff Fara resulted in too many small changes to my views over time to catalogue in a short preface. This work has been presented in far too many locations for me to remember who else has contributed to its development. I apologize to those whose contributions I have failed to acknowledge.

The chapters in this volume were written while I was on the faculty of Cornell University and the University of Michigan, Ann Arbor. I am deeply grateful to my colleagues at both institutions for the intellectual community they helped foster. I wrote the introduction and the postscript to this volume while on sabbatical from my current position at Rutgers University. I am grateful to my colleagues for their contributions to my thinking, and to the institution for its support. “Context and Logical Form” originally appeared in *Linguistics and Philosophy*, 23/4 (2000): 391–434, “On Quantifier Domain Restriction” originally appeared in *Mind and Language*, 15/2&3 (2000): 219–61, and “Making it Articulated” also originally appeared in *Mind and Language*, 17/1&2 (2002): 149–68. I am grateful to the editors of these journals for permission to reprint these articles here. Two chapters in this book originally appeared in collections edited by Gerhard Preyer and George Peter. “Nominal Restriction” originally appeared in *Logical Form and Language* (Oxford: Clarendon Press, 2002), 365–88; and “Semantics



in Context” appeared in *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Oxford University Press, 2005), 221–53. “Semantics, Pragmatics, and the Role of Semantic Content” also originally appeared in a collection by the same publisher, *Semantics versus Pragmatics* (Oxford: Clarendon Press, 2005), but this time edited by Zoltán Gendler Szabó. Thanks to Clarendon Press for permission to reprint these articles. My review of François Recanati’s book originally appeared in the *Notre Dame Philosophical Reviews*, online. Thanks also to Nadiah Al-Ammar at Oxford University Press, who shepherded this manuscript through the process of publication, to Sally McCann for her copy-editing, to David Carles for proofreading, and to Karen Lewis for compiling the index. Thanks especially to Allan Espiritu of Rutgers University for such a creative cover design.

Linguists and philosophers are not the only people who aided me in the writing of these papers. My mother deserves great thanks for her unwavering support of my endeavors. My wife, Njeri, has watched me struggle through every idea in this book, and has provided much needed perspective and refuge. Of all the people in my life, she most obviously deserves a book dedicated to her. I hope this is but the first of many.

# CONTENTS

|  |     |
|--|-----|
| Introduction   | 1   |
| 1. Context and Logical Form  | 30  |
| 2. On Quantifier Domain Restriction (with Zoltán Gendler Szabó)                      | 69  |
| 3. Nominal Restriction   | 111 |
| 4. Semantics, Pragmatics, and the Role of Semantic Content<br>(with Jeffrey C. King) | 133 |
| 5. Making it Articulated   | 182 |
| 6. Semantics in Context  | 201 |
| 7. Review of François Recanati's <i>Literal Meaning</i>                              | 231 |
| Postscript   | 248 |
| <i>Index</i>   | 261 |

*This page intentionally left blank*

# Introduction

A conversation involves acquiring and conveying information about the world, immediately and without much conscious reflection. Our linguistic capacity thereby enables us to engage in complex cooperative activities that require rapid information flow between large numbers of people, such as building bridges and superconductors (though unfortunately also waging war). We share our ability to perceive visually with many other species. In contrast, no other species has our advanced linguistic capacity. Any investigation into what makes humans distinctive as a species begins and ends with language-use. The inquiry into the nature of our linguistic capacities is therefore at the very center of the human sciences.

Perhaps because it is so unique among the species, progress in understanding our capacity for gaining and conveying information about the world linguistically has been slower than progress in understanding (say) our perceptual capacities. We have made many inroads in grammar and phonology. But a complete account of the relation between an utterance and the information about the world it conveys is still probably no closer than a complete account of the biology of the species. A complex phenomenon needs to be broken up into many components to have any hope of ultimately being explained. The phenomenon of language-use is no different in this regard. My proposal is a contribution to only a small part of an account of the relation between utterances and the information about the world they convey. For example, there are very difficult questions about what the meanings of words are, and even more difficult questions about the relations of words to the meanings they have. Is the meaning of the word “Mars” the planet Mars? Or is it a rule of use for employing the word “Mars” correctly? If the meaning of the word “Mars” is the planet Mars, how did this relation between a word and a planet come into being? My project in this book only indirectly concerns these difficult questions. I will therefore presuppose without question some standard answers. For instance, because we learn something about the planet Mars from an utterance of “Mars is the nearest planet to Earth”, I will assume that the semantic content of “Mars” is the planet Mars. But no defense of this view of meanings will be given here. Furthermore, I have nothing to say, in the chapters that follow, about the thorny question of the relation between individual words and the

## 2 / Language in Context

things in the world that are their meanings. For example, I will have nothing to say about whether the relation between the words “Mars” and the planet Mars is a causal relation, or even a naturalistically acceptable one. I simply assume, for purposes of my inquiry, that the word “Mars” does refer to the planet Mars. My project begins where questions about standing linguistic meaning leave off.

Once we know the standing linguistic meanings of all the words in a sentence, and the grammar of the sentence, there is still a gap between that knowledge and the information about the world conveyed by an utterance of that sentence. For example, an utterance of the sentence “Jack is tall” may convey the information that Jack is tall for a fourth-grade student. Nothing in the standing linguistic meanings of “Jack” or “tall” seems to determine this information. The chapters in this book consist of a defense of one hypothesis about the gap between the standing linguistic meaning of the words in a sentence and the information conveyed by an utterance of it. The defense is not complete, and to even make the sort of hypothesis I do in this work is no doubt premature. But concrete, developed proposals are needed at this stage of inquiry. Even if they turn out to be false, we will at least learn something from their failures.

I have said that my project only indirectly bears on foundational questions such as whether meanings are best conceived of as things in the world (such as the planet Mars) or rules of use. It should be clear that it does not *directly* bear on this question, since I take it for granted that the standing linguistic meanings of non-context sensitive words are often objects, properties, or events in the world. For example, I take it for granted that the standing linguistic meaning of “is taller than six feet” is the property of being taller than six feet, which some people have and others lack. But my project does *indirectly* bear on the question whether we can appropriately explicate meaning in terms of reference to objects and properties in the world. Ordinary language philosophers have argued strenuously that properties such as truth and reference do not apply to linguistic expressions but are rather properties of what people *do* with linguistic expressions. An object is not the meaning of a singular term, and a proposition (an entity capable of truth or falsity) is not the meaning of a sentence. Rather, it is a *use* of a singular term by a person that refers, and it is an *utterance* of a sentence that expresses a proposition; one cannot speak of a *term* having reference, or a *sentence* having a truth-value, even derivatively. In short: words do not refer, people do. If semantic properties such as referring to an object or being true are not properties of linguistic expressions, then giving an account of linguistic meaning in terms of reference and truth is fundamentally misguided. Philosophers such as Carnap and Tarski were right to focus their attention on formal languages, because the kind of account of meaning they were trying to give (in terms of reference and truth) was inapplicable to natural languages.

Since referring to an object is not a property of a linguistic expression, and linguistic meanings are properties of linguistic expressions, ordinary language philosophers sought an alternative account of linguistic meaning. According to it, the linguistic meaning of an expression is a *rule for its proper use*. The motivation ordinary language philosophers gave for the view that linguistic meaning was best explicated in terms of rules of use centrally involved *context-sensitivity*. The fact that one and the same sentence could be used to convey quite different information in different contexts of use seemed incompatible with giving a theory of meaning in terms of reference and truth. First, ordinary language philosophers argued that the existence of context-sensitive words posed a problem for accounts of word meaning in terms of reference. As Peter Strawson writes in “On Referring”:

If someone asks me the meaning of the expression “this” . . . I do not hand him the object I have just used the expression to refer to, adding at the same time that the meaning of the word changes every time it is used. Nor do I hand him all the objects it ever has been, or might be, used to refer to. I explain and illustrate the conventions governing the use of the expression. This *is* giving the meaning of the expression.

Secondly, the gap between the kinds of accounts of meaning in terms of reference and truth and the intuitive content of utterances seemed large enough as to vitiate the enterprise. For example, for a long time, the only model of the meaning of sentential connectives such as “if . . . then” and “or” was in terms of truth-functions. Yet the ordinary language meaning (and context-sensitive nature) of these sentential connectives was not adequately captured by truth-tables.

There were two kinds of responses to ordinary language philosophers on behalf of philosophers who thought of meaning as best explained in terms of reference and truth. First, in his extraordinarily influential paper “Logic and Conversation” (Grice, 1989a), the philosopher Paul Grice set out to defend the truth-table analysis of the meaning of the natural language sentential connectives from the ordinary language onslaught. Ordinary language philosophers often did not distinguish the *truth* of an utterance from the *acceptability* of that utterance. The key to Grice’s defense of the truth-table analysis of the meanings of “and”, “or”, and “if . . . not” is that these notions can (and often do) come apart. A given utterance can be true, even though uttering it is not acceptable, because it violates conversational norms. In explaining this distinction, Grice provided the foundations for a theory of conversational norms. The theory Grice gives clearly explains how an utterance may be true, though unacceptable as an assertion due to specific facts about the conversation and its participants. Grice then used the distinction between the truth of a statement and its conversational acceptability in a defense of the thesis that the connectives of propositional logic were correct explications of their natural language counterparts. More specifically, Grice assumed that the natural language logical particles have the truth-table meanings of

#### 4 / Language in Context

their logical counterparts, and argued that features of the uses of these expressions that are not explicable by the truth-tables are due to facts about the norms governing conversation, rather than the meanings of the words.

Though most philosophers no longer accept many aspects of Grice's positive account of the meaning of natural language sentential connectives, everyone accepts that some gaps between what an utterance says and what it conveys are due to general conversational norms, rather than facts about meaning. Thus, appeal to such gaps cannot be used to undermine an account of meaning in terms of reference and truth. The fact that an instance of "P or Q" is only felicitous when the person uttering it does not know that p and does not know that q cannot be used as a premise in an argument for the conclusion that "or" does not refer to a truth-function. For this fact about "or" is due to the general conversational norm that one should convey as much information as one knows. So if one knows that p, one should assert that p, rather than asserting the weaker proposition that either p or q. If Hannah knows that John was at the party, it would be a violation of the Grice's *maxim of quantity* for her to assert that either John was at the party or he was at home. She would not be being maximally informative by asserting the disjunctive statement, and hence would be violating the maxim of quantity. The oddity of Hannah's utterance is fully consistent with, and indeed explained by, the fact that "or" refers to a certain truth-function.

The second kind of response to the ordinary language philosopher emerged from the increasing sophistication of semantic theories. Recall that Strawson ridiculed the notion that one could use the notion of reference to give an account of the meaning of a context-sensitive word, such as "this". But the philosopher Richard Montague, and his students Hans Kamp and (in particular) David Kaplan, developed semantic theories where the truth-conditions of a sentence were given *relative to a context of use*. For example, on the semantic theory developed by David Kaplan in his seminal paper "Demonstratives" (Kaplan, 1989), it makes perfect sense to speak of a context-sensitive singular term having reference, albeit relative to a context of use. The context-independent meaning of a context-sensitive singular term, such as the English first-person pronoun "I", is modeled, on such theories, as a *function* from contexts to persons. For example, the context-independent meaning of "I" is a function that takes a context, and yields as a referent the person speaking in that context. So when Bill Clinton utters "I am tired", he expresses the proposition that Bill Clinton is tired, because the context-independent meaning of "I", as applied to a context in which Bill Clinton is speaking, yields Bill Clinton as value. On this approach, it makes perfect sense to attribute reference and truth to expression types, once contextual relativity is factored into the semantic theory. Whereas the notion of a rule of use is vague and mystical, Kaplan's notion of the standing linguistic meaning, or (as he calls it), the *character* of an expression, is not only clear, but set theoretically explicable in terms of fundamental semantic notions; the character of an expression

is a function from a context to the reference of that expression in that context. Far from context-sensitivity being an impediment to giving a proper account of linguistic meaning in terms of reference and truth, appeal to these semantic notions allows us to give a considerably more explicit characterization of linguistic meaning than the ordinary language philosophers were capable of providing.

So Grice, on the one hand, and Montague and Kaplan, on the other, together provided a systematic defense of the view that meaning is best explicated in terms of reference, rather than rules of use. If Montague and Kaplan are right, it makes perfect sense to attribute reference to word-types, albeit relative to a context. But these philosophers did not thereby defeat the ordinary language argument for the conclusion that truth and falsity cannot be properties of sentences, but only of acts of uttering those sentences, and the propositions such acts express. For, even given an account of the standing meaning of context-sensitive words in terms of functions from contexts to referents, there is still a large gap between the content a sentence appears to have, relative to a context, and the information it intuitively conveys. Furthermore, this gap cannot be explained away by general conversational norms.

For example, consider again the sentence “Jack is tall”. Uses of this sentence can express different propositions, relative to different contexts of use. One use of the sentence can express the proposition that Jack is tall for a fourth-grader. Another use of the sentence can express the proposition that Jack is tall for a basketball player. But general conversational norms surely do not explain the gap between the linguistically determined content of “Jack is tall”, relative to a context of use, and the proposition it intuitively is used to express. Indeed, the linguistically determined content expressed by “Jack is tall”, even relative to a context, does not seem susceptible of truth and falsity at all. It is only when we add the additional contextually provided information—tall relative to *what*, that we are capable of assessing it for truth and falsity. This suggests that the ordinary language philosophers were correct in maintaining that properties such as truth and falsity cannot be said to belong to sentences, even relative to a context. Perhaps, then, meanings are after all better thought of as rules for expressing contents in context, i.e., rules of use.

My project in the chapters in this book is to respond to this last remaining vestige of the ordinary language philosopher’s challenge. If my proposal is correct, there is no gap between the linguistically determined content of a sentence, relative to a context, and the proposition it intuitively seems to express. General conversational norms, plus an account of meaning in terms of reference and truth (a *truth-conditional semantics*), explains the gap between grammar and what is conveyed. I now turn to a more detailed explanation of my project.

Suppose that Hannah utters a sentence in a conversation, say the sentence “Bill Clinton lived in Arkansas”. In uttering this sentence, Hannah imparts certain information about the world. Intuitively, the information that Hannah imparts to



## 6 / Language in Context

her interlocutors is that a certain person, Bill Clinton, lived in Arkansas. As native speakers of the language, we know that this is what has to be true about the world, in order for Hannah's utterance of the sentence "Bill Clinton lived in Arkansas" to be true. Accordingly, we may call this information the truth-conditions of Hannah's utterance. Ordinary speakers have fairly robust intuitions about the truth-conditions of an utterance. For example, consider an utterance of the sentence "Aristotle was a great philosopher". We all have the intuition that what this utterance says would be false, relative to a situation in which Aristotle had died in infancy. The conditions for the truth of what is said by an utterance of "Aristotle was a great philosopher" are simply not met relative to a situation in which Aristotle died in infancy. Similarly, consider an utterance of "John is tired and Sue is tired". Is what is said by this utterance true relative to a situation in which only John is tired? Clearly, it is not. The truth-conditions of an utterance of "John is tired and Sue is tired" require both John and Sue to be tired. As native speakers of the language, we have robust intuitions about the truth and falsity of what is said by an utterance of English relative to different possible situations.

Speakers have robust intuitions about the grammaticality and ungrammaticality of sentences of their native languages. This fact has been used by linguists to help construct theories of the grammars of natural languages; the grammaticality intuitions we have about our native tongues form the evidential basis for the hypotheses of syntactic theory. The fact that native speakers of a language have robust intuitions about the truth and falsity of what is said by utterances of sentences of that language has been similarly exploited by theorists across a number of disciplines interested in developing theories of meaning for natural languages. In philosophy of language, theorists have exploited these intuitions to argue for many conclusions about meaning. For example, Saul Kripke has exploited such intuitions to argue that names are not disguised definite descriptions, as Bertrand Russell had claimed. In semantics, linguists have exploited these intuitions to develop accounts of the meaning of a wide range of constructions. In short, intuitions about the truth and falsity of what is said by utterances of sentences have formed the data by which theorists have tested their hypotheses about meaning. There is no other obvious source of native speaker intuitions that are related to meaning. So if we did not have robust intuitions about the truth-conditions of our utterances, it would not be clear how to test such hypotheses; there would be no firm basis on which to construct a theory of meaning.

Native speakers possess intuitions about the truth and falsity of what is said by utterances of sentences of their native languages, relative to various possible situations. This fact has been exploited to construct accounts of the meanings of particular constructions in natural language, and indeed it is the only viable basis for constructing an account of meaning for a natural language. But what is the relation

between an utterance and its intuitive truth-conditions? Or, more specifically, how does the utterance come to be associated with those truth-conditions?

It is natural to explain how a specific utterance comes to be associated with the truth-conditions it seems to have (to a rational speaker in good communicative conditions) by adverting to the meanings of the constituent words in the sentence uttered, together with their occurrence in the grammatical structure of that sentence. Perhaps an utterance has the truth-conditions it seems to have because the words in the sentence being uttered have a certain meaning, and they are combined in a certain way in the sentence. When so combined, these meanings determine the truth-conditions an utterance of that sentence seems to have. An utterance of the sentence “Bill Clinton lived in Arkansas” has the intuitive truth-conditions it does because the meanings in English of “Bill Clinton”, “lived”, “in”, and “Arkansas”, when combined in accord with the grammar of the sentence, uniquely determine these truth-conditions. An utterance of the sentence “Bill saw Hannah” has distinct truth-conditions from an utterance of “Hannah saw Bill” because the two sentences, while containing the same words, contain them in distinct grammatical positions.

No doubt, the fact that the words that occur in a sentence have the meaning and grammatical position they do plays a substantial role in any explanation of why an utterance of that sentence has the intuitive truth-conditions it does. But this cannot be the entire explanation of why utterances have the intuitive truth-conditions they do. For it is uncontroversial that these features of a sentence underdetermine the truth-conditions of an utterance of that sentence. For example, the English sentence “I am human” has a single grammatical structure (i.e., it is not structurally ambiguous) and it is composed of words, each of which has a single, unvarying linguistic meaning. But different utterances of this sentence intuitively have different truth-conditions. If Bill Clinton utters this sentence, the intuitive truth-conditions of his utterance are that Bill Clinton is human, whereas if I utter this sentence, the intuitive truth-conditions of my utterance are that Jason Stanley is human. So it cannot be that the linguistic meanings of the words used in a sentence together with their grammatical positions provide a complete explanation of why utterances have the intuitive truth-conditions they do.

The fact that the words in the sentence being uttered have the linguistic meanings and grammatical position they do therefore is part of the explanation of why that utterance has the intuitive truth-conditions it does, but it is not a full explanation. Facts about the extra-linguistic context in which the sentence is uttered help determine its intuitive truth-conditions in that context. The investigations in this book are devoted to explaining how extra-linguistic context, together with linguistic meaning and grammatical structure, gives rise to content.

A number of theorists are pessimistic about the project of explaining how extra-linguistic context interacts with linguistic meaning and structure to yield the

intuitive truth-conditions of an utterance; they hold that beyond linguistic meaning and structure, there is nothing systematic to be said about the relation between an utterance of a sentence and its intuitive truth-conditions. On this view, there are ever so many reasons why utterances have the intuitive truth-conditions they do, and nothing of great import to say about what unifies these disparate explanations. At best, what a theorist can do is say what the linguistic meaning of the words pronounced in a sentence are, and describe the grammatical structure of the sentence. Furthermore, so doing does not allow us to determine what the intuitive truth-conditions of utterances of that sentence would be. Call this broad range of views, the *pessimistic view*. In contrast, one who adopts an optimistic attitude towards the study of the relations between utterances and their intuitive truth-conditions holds that there is an illuminating uniform description of the relation between utterances and their intuitive truth-conditions, beyond just a specification of context-independent linguistic meaning and grammatical structure.

On the pessimistic view, there is stability to word meaning and the significance of the syntactic structure of sentences. But in general there is no systematic way of going from the meanings of the words in a sentence and its syntactic structure to the intuitive truth-conditions of its various utterances. Traditionally, one motivates this position by arguing that the unsystematic way in which the context of discourses affects the intuitive truth-conditions of the utterances that occur in them undermines the search for an illuminating general description of the mapping between utterances and their intuitive truth-conditions. According to Noam Chomsky, for example, there is no scientific study (i.e., no interesting systematic account) of the relation between utterances of sentences and the truth-conditions they intuitively possess. There is still, according to Chomsky, a science of meaning. But such a science cannot investigate the relation between language and the world, since there is no systematic account of this relation (or there is no such relation). Instead, it must involve relations between our sentences and their “I-Meanings”, which, on Chomsky’s view, are mental entities of some kind. Chomsky’s position has been ably defended in recent years by philosophers such as Paul Pietroski.

If any version of the pessimistic view were correct, significant facts about linguistic communication would be inexplicable. From an utterance of a sentence, one gains information about the world. If Hannah utters to Esther the English sentence “There is some chocolate in the kitchen”, and if Esther wants chocolate, she will go to the kitchen. Even if Esther has never heard that particular sentence before, if she speaks English, she will understand what information it conveys about her physical environment. Though the pessimist gives us an account of the relation between Hannah’s utterance and its I-meaning, what spurs us to action is the information the utterance conveys about the world. If the pessimist is correct, there is nothing illuminating and systematic to say about *how* language conveys information about the world. The pessimistic view is difficult to accept because language-users (even at very young ages)

smoothly grasp information about the world from sentences they have never previously encountered. Furthermore, they do so given only knowledge of the meanings of a (relatively) small number of individual words and their modes of combination into sentences. If there is no systematic way of proceeding from knowledge of the extra-linguistic context and knowledge of the meaning of individual words and modes of combination into sentences to a grasp of the information about the world that is conveyed by an utterance of a sentence, it is mysterious how language-users could so smoothly move from linguistic comprehension to action. In short, if the pessimistic view were correct, the connection between speech and action would be inexplicable.

Furthermore, if the pessimistic view were correct, it would not be clear how to construct a theory of meaning. Native speakers have intuitions about the truth and falsity of what is said by an utterance, relative to various possible situations. But native speakers do not have intuitions about meaning. Even dictionaries are a poor source for discovering the “I-meanings” postulated by theorists such as Chomsky. In short, theorists of meaning who do not think there are systematic relations between the intuitions native speakers have about the truth and falsity of utterances and the meanings of words and sentences have stripped themselves of any plausible evidential basis for their hypotheses.

If the pessimistic view is incorrect, then there must be some way to account, in a systematic manner, for the apparently significantly diverse ways in which linguistic interpretation is conditioned by context. That is, it is incumbent upon those who reject the pessimistic approach to produce a systematic account of the data to which advocates of pessimism have drawn our attention over the years. Meeting this challenge is complicated by the different strategies available for pursuing a systematic account of linguistic interpretation.

One possible way to pursue a systematic strategy for explaining the relation between utterances and the truth-conditions they intuitively possess is to attempt a systematic explanation of rational communicative action generally. On this model, advocated for example by *relevance theorists* such as Dan Sperber and Deirdre Wilson (1986) and Robyn Carston (2002), there is a systematic account of interpretation of the mental states of others. It is our ability to make reasonable inferences from the observable behavior and characteristics of our interlocutors to their mental states that ultimately explains our success in acquiring information from their utterances. On this model, linguistic behavior is no different in kind from the other cues (winks and nods, half-smiles) our interlocutors give about their inner narratives. The way we draw out information from others is by interpreting, in the first instance, their mental states; their utterances are just evidence, no different in kind from the evidence provided by the kind of clothing they tend to wear.

Advocates of this strategy do think that the meanings of the words in a sentence help us in deciding what the speaker intends to communicate by her utterance of that

sentence. But they also think that the linguistically guided information provided by the sentence and its grammatical structure is quite minimal. In interpreting others, we employ non-linguistically guided reasoning about their mental states. This allows us to enrich the linguistically guided information we glean from the sentence they use with other information, via a process advocates of relevance theoretic pragmatics call *free enrichment*, and advocates of alternative pragmatic approaches, such as Kent Bach (1994), call *implicature*. The intuitions speakers have about the truth or falsity of what is said by an utterance is thoroughly influenced by these non-linguistically guided factors.

It is not clear that this strategy can meet the pessimists' challenge. On this view, there is no direct systematic relationship between utterances and the truth-conditions they seem to have. Instead, the systematic relationship is between two interlocutors interpreting one another's mental states. Indeed, there may be no way, on this view, of setting up any kind of systematic relation between utterances and the truth-conditions they seem to have. More worryingly, such a view does not seem to account for the smoothness and rapidity of linguistic communication. When someone hears a sentence containing context-sensitive terms such as "that", "he", or "she" under good communicative conditions, they rapidly grasp the specific truth-conditions of the utterance. Of course, to decode the contribution of the context-sensitive terms, a hearer must appeal to facts about the discourse context, including salient facts about the mental states of her interlocutors. But the fact that we do so quite rapidly suggests that our search for the facts to interpretation is significantly constrained by linguistic meaning, that is, is under linguistic control.

Many philosophers of language, impressed by the challenge of reflecting the pervasive dependence of utterance content on context, have argued that a theory of meaning for a language should not incorporate much context-dependence. According to this *minimalist* approach, the semantic content of a sentence relative to a context bears only a loose relation to the intuitive truth-conditions of an utterance of that sentence in that context; our intuitions about the latter are the result of both the semantic content and myriad unsystematic pragmatic effects (e.g. Borg (2004), Cappelen and Lepore (2004)). For example, suppose a teacher utters the sentence, "Everyone has passed the exam" to her class. Minimal semanticists maintain that the semantic content of her sentence, in that context, is the proposition that everyone in the universe passed the one and only one exam in the universe. The proposition we intuitively feel is expressed by her utterance is the result of (say) recognizing that this is an absurd proposition to assert, and pragmatically adjusting to arrive at the proposition that everyone in the class passed the exam given in that class. The minimal semanticist postulates informationally impoverished semantic contents of sentences, to avoid admitting context-sensitivity in the theory of meaning.

The philosophers who espouse minimal semantics hope to save systematic semantics from the challenges posed by pervasive context-sensitivity. However, speaker intuitions about the truth and falsity of sentences relative to contexts are the evidential basis for hypotheses about meaning. Minimal semantics severs the data of the theory of meaning from its hypotheses, rendering the semantic project a tapestry of idle speculation. Perhaps for this reason, the project of minimal semantics has few serious adherents in linguistics. Minimal semanticists have had little or nothing to say about how the theory of meaning should be carried out in the absence of any data for its claims.

The popularity of the minimal semanticist's perspective in philosophy of language is due to the fact that philosophers tend to detect some threat from context-sensitivity to the project of the theory of meaning. However, the nature of this putative threat is far from clear. It is true that some of the progenitors of the analytic tradition in philosophy were suspicious of the use of natural language in reasoning, because of context-sensitivity and ambiguity. But this is because their chief concern was not the theory of meaning, but rather the project of justifying arithmetic. To justify arithmetic, they sought to derive the arithmetical axioms from logical axioms within a system that was plausibly a formalization of pure logic. Carrying out mathematical proofs in a context-sensitive language adds additional risk of fallacy. If a sentence *S* is context-sensitive, then it can express different propositions, relative to different contexts of use. If the context-sensitive sentence *S* occurs multiple times in a proof, say once as a premise and another time as an antecedent of a conditional, then it may express different propositions in these different occurrences. If so, then while the syntactic form of the inference may be valid (in this case, an instance of *modus ponens*), the inference itself will not be valid. So a context-sensitive language is not an ideal language in which to carry out mathematical proofs; such a language introduces novel possibilities of error. It is for this reason that the progenitors of analytic philosophy, in their attempt to devise an ideal formalism for conducting mathematical proofs, eschewed context-sensitivity. I suspect that much of the aversion philosophers have to context-sensitivity in a theory of meaning is a remnant of a legitimate aversion to context-sensitivity by our philosophical ancestors, given the aims of their projects. But natural language semantics is a very different project than logicism; it is an error to let the concerns of the latter affect us when engaged in the former.

There are additional disadvantages to employing a language rife with context-sensitive constructions. For example, as Timothy Williamson (2005, 100) notes, one cost of pervasive context-sensitivity is that it complicates "the preservation of information in memory and its transmission by testimony" (see also Hawthorne (2004, 109–10)). If adjectives such as "tall" introduce a kind of context-sensitivity, say to a comparison class or a degree of height, then in order to transmit the information I convey by an utterance of "Jack is tall", you must identify the exact comparison class

or degree of height I have in mind. Similarly, in order for me to remember what I believed when I believed that Jack is tall I must know exactly what comparison class I had in mind.

So, context-sensitivity complicates preservation of information. It makes communication somewhat more difficult in general, if we think of successful communication as conveyance from speaker to hearer of the very same proposition. If it is a conversational maxim, as some have suggested, that “speakers ought, in general, to assume that their addressees have whatever information is necessary to determine what they are saying” (Stalnaker (1987, 110)), then pervasive context-sensitivity will also make adherence to conversational norms more difficult. Would it be better, then, to reject the appearance of context-sensitivity in natural language, and construct a theory of meaning that ignores it? Clearly it would not be. As we have seen, arguments presented in a context-sensitive language face an additional possibility of error. In this sense, giving arguments in a context-sensitive language is risky. But if context-sensitivity is already present in the language, not recognizing it is considerably riskier. Arguments may seem valid, which in fact are not. Similarly, if we do convey information with the use of pervasively context-sensitive constructions, disregarding context-sensitivity is not a way to preserve the transmission of knowledge by testimony. What would thereby be preserved would not be the information that was in fact originally conveyed by the original utterance, but some irrelevant informationally impoverished monstrosity; contents that are sufficiently impoverished to be faithfully transmitted are not sufficiently interesting to assert. Minimalist semantics is an attempt to retain a fanciful philosopher’s illusion, at the cost of exposing us to inferential risk and emphasizing informational absurdity.

Finally, minimal semantics helps not a whit in explaining the problems with which we started. Our problem was to explain how humans rapidly go from sentences they have never before encountered to information about the world. The minimal semanticist can explain how humans rapidly go from sentences they have never before encountered to minimal content. For example, the minimal semanticist can explain how humans rapidly go from hearing the sentence “Everyone passed the exam” to the minimal content that everyone in the universe passed the only exam in the universe. But that is where her explanation ceases. And there are two problems with this project. First, it is not clear that speakers ever do rapidly go from hearing an utterance of the sentence “Everyone passed the exam” to the minimal content that everyone in the universe passed the only exam in the universe; indeed, it is not clear that speakers ever entertain minimal contents at all. The minimal semanticist is therefore providing an explanation for phenomena that never in fact occur. Secondly, even if the minimal semanticist could make it plausible that language users do tacitly entertain minimal contents, they would still be addressing a distinct problem from the one with which we started, which was to explain how language-users rapidly

go from sentences they have never before encountered to information about the world that spurs them to action. Minimal contents are not the kinds of things that spur language-users to action; grasping the proposition that everyone in the universe passed the only exam in the universe is not particularly useful. To provide the missing explanation, the minimal semanticist would have to spend the bulk of her time developing a sweeping pragmatic account of the relation between sentences and their intuitive truth-conditions, of the sort that relevance theorists or other pragmatic theorists are trying to develop. In other words, even if language-users did tacitly entertain minimal contents, minimal semantics would still end up being nothing but a branch of, for example, relevance theoretic pragmatics.

So how do we communicate, if most of our sentences have multiple context-sensitive elements? It may be, as Richard Heck (2002) has argued, that we have been mistakenly seduced into accepting “the naïve conception of communication”, according to which successful communication requires grasp of the same proposition expressed by an utterance of a sentence. Perhaps, as Heck and others have suggested (e.g., Bezuidenhout (1997)), successful communication requires only grasp of a *sufficiently similar* proposition to the one expressed by the utterance. Cappelen and Lepore (2006) rightly protest that it is particularly difficult to give a reductive characterization of the conditions under which two propositions are sufficiently similar in a given context. But perhaps we do not need such a characterization to reject the naïve conception.

I am in fact attracted to a quite different response to the problem than the one entertained by Heck, Bezuidenhout, and others. I am skeptical about the grounds philosophers have given to think that it is very difficult to grasp the proposition expressed by someone else’s utterance of a context-sensitive sentence. First, it is tempting to think that successful communication requires the interlocutor to think of the objects and properties that constitute the proposition in the same way as the person who utters the sentence thinks of those objects and properties. Second, it is tempting to think that grasping a property requires being able to distinguish that property from very similar ones. Third, it is tempting to think that one grasps the proposition expressed by someone else’s utterance only if, had one uttered the same sentence in a similar context, one would thereby express the same proposition. But all of these temptations should be resisted.

Suppose someone utters a sentence containing a quantified noun phrase. They thereby express a proposition containing a property that is the domain of that quantified noun phrase. For an interlocutor to grasp that proposition, it is not necessary to think of that property in the same way as the way the person who utters that proposition. For example, suppose I utter the sentence, pointing at John, “He is tired”. I may think of John somewhat differently than you think of John. But I can still grasp the singular proposition about John that you express. A similar point holds



for properties. It is no doubt a difficult matter to state the conditions under which someone has a *de re* thought about an object or a property, and the conditions for *de re* acquaintance with an object or property might themselves be situation dependent. But it is a difficult matter that has nothing specifically to do with context-sensitivity in the philosophy of language.

The availability of this response depends, of course, on the conception of proposition that one employs. In giving this response, I have characterized the contents of utterances as Russellian propositions, which are structured meanings that contain objects and properties. If one instead employs a semantic theory with *Fregean* propositions, as Heck advocates, then this response is more difficult to give. Fregean propositions contain ways of thinking of objects and properties. To grasp a proposition containing ways of thinking of objects and properties, it seems one must grasp the specific ways of thinking that constitute the proposition. Explaining how interlocutors can “share content” (to use the apt expression of Cappelen and Lepore (2006)) is, as one might have suspected, more difficult to solve if meanings are individuated in terms of ways of thinking, rather than in terms of objects and properties.

The second thought that leads philosophers to think that grasping the content expressed by someone else’s utterance of a context-sensitive sentence is difficult is that an interlocutor is often unable to distinguish the property intended by the speaker as (for example) the intended domain of quantification from other similar properties. For example, suppose John utters the sentence “Everyone is a philosopher”, meaning to express the proposition that everyone in the room is a philosopher. But suppose everyone in the room is sitting on the right side of the room. One might think that unless John’s audience knows that he intends the property of being in the room *rather than* the property of being on the right side of the room, they do not grasp the proposition he expresses. In short, the thought is that a person does not grasp a proposition expressed by a sentence unless they can distinguish the properties in that proposition from all properties that are sufficiently similar, if they were presented with all the possible alternatives at the same time.

But it is unclear, to say the least, why we should accept the epistemic requirement on grasp of the constituents of a proposition that motivates this line of thought. If someone utters the sentence “Water is wet”, surely I can grasp the proposition thereby expressed, even if I cannot distinguish water from very similar substances. Grasping a proposition surely does not require meeting such a demanding epistemic requirement. If it did, we would have trouble explaining how we grasp the propositions expressed by many eternal (non-context-sensitive) sentences.

The third tempting thought is that one grasps the proposition expressed by someone else’s utterance of a sentence *S* only if, were one to utter the same sentence in the same or similar context (obviously the speaker of the context would be different),

one would express the same proposition. For example, Cappelen and Lepore (2006, section 3.2) write:

Consider two sailors on the ship, Popeye and Bluto. After the sad departure, Popeye observes “That was a nice occasion. Every sailor waved to every sailor”. Immediately afterwards, Bluto concurs, “That’s right. Every sailor waved to every sailor”. In such a circumstance the following is often obvious: we treat these utterances as expressing agreement. Popeye and Bluto agree that every sailor waved to every sailor . . . But if Stanley and Szabo [whose theory of quantifier domain restriction is under dispute] were right about the semantics of quantifiers, their concurrence would be a minor miracle.

There are several thoughts contained in this passage. The first is that Bluto grasped what Popeye said only if Bluto’s utterance of the same sentence expressed the same proposition as Popeye’s previous utterance. The second is that the semantics of ascriptions of “agree” requires sameness of content. I will not challenge the second of these thoughts here. But the first seems false.

Suppose speaker intentions determine the value of quantifier domain indices, and suppose, for the sake of argument, that Bluto intends to refer to slightly different domains for his two uses of “every sailor” than Popeye did. Then their two utterances express different propositions. But it simply does not follow that Bluto did not grasp the proposition expressed by Popeye. Suppose Popeye expressed the proposition that every sailor on the ship waved to every sailor on the shore. So, the domain for Popeye’s first use of “every sailor” is the property of being on the ship, and the domain for Popeye’s second use of “every sailor” is the property of being on the shore. Bluto could grasp this proposition in this context by thinking of the property of being on the ship (which is the domain for the first quantified noun phrase used by Popeye) as the unique ship-related property intended by Popeye, and he could grasp the property of being on the shore as the unique shore-related property intended by Popeye. In this context, this would suffice to give him a *de re* grasp of the quantifier domain properties contained in the proposition expressed by Popeye’s utterance of “every sailor waved to every sailor”. So, whether or not Bluto succeeds in expressing the same proposition by his subsequent utterance of the same sentence, he can grasp what is said by Popeye.<sup>1</sup>

So I do not see the case for thinking that ubiquitous semantic context-dependence threatens either the possibility or the systematicity of interpretation. I therefore

<sup>1</sup> These thoughts can quite obviously be brought to bear on the conclusions that have been drawn in the literature on indicative conditionals from the example of “Sly Pete”, introduced in Gibbard (1981). Sly Pete examples and variants thereof have been taken to undermine the thesis that indicative conditionals have truth-conditions, on the grounds that if they did have truth-conditions, their truth-conditions would vary too much as a function of context. I think similar considerations would show that many more kinds of sentences in our language fail to have truth-conditions, including a wide variety of sentences containing modal expressions.

have sought to develop a systematic account of the apparently unsystematic ways in which the context of discourse may affect the truth-conditions of the utterances that constitute it, one that does not proceed via an appeal to a systematic general theory of rational communicative action.

The account I have been developing and defending is only one possible systematic account of the relation between utterances and the truth-conditions they seem to have. But it is perhaps the most straightforward such account. Consider a sentence *S* as uttered in a context *c*. This utterance of *S* has the truth-conditions it seems to have because these truth-conditions are what results from taking the words contained in *S*, considering what their content is relative to the context *c*, and combining these contents in accord with the syntactic structure of *S*.

So, for example, consider the sentence “I am human”, uttered by Hannah. Intuitively, Hannah’s utterance is true if and only if Hannah is human. These are the truth-conditions of Hannah’s utterance, on my view, because they come from taking the *content* of “I” relative to the context in question (which is Hannah), and the content of the predicate “am human” relative to that context (which is the property of being human). Combining these in accord with predicational structure yields the truth-condition that Hannah is human. Similarly, consider the sentence, “She is tired now”, uttered by Hannah at 4 p.m. on 12 July 2003 (and suppose Hannah is pointing at Sara). Intuitively, Hannah’s utterance is true if and only if Sara is tired at 4 p.m. on 12 July 2003. These are the truth-conditions of Hannah’s utterance, on my view, because the content of Hannah’s use of “now” is the time of utterance (4 p.m. on 12 July 2003), the content of Hannah’s use of “She” is Sara, and combining this with the content of “is tired” yields those truth-conditions.

We have the ability to gain information about the world from the utterances of others, given only knowledge of the meaning of words and the significance of combining words into more complex expressions. If the explanation of the relation between utterances and their intuitive truth-conditions I have just provided is correct, this ability is not mysterious. It is explained by our grasp of the *contents* (rather than just the linguistic meanings) of individual words and the significance of combining them, for, if I am right, together this yields the intuitive truth-conditions of an utterance.

Of course, some words make different contributions in different contexts to the intuitive truth-conditions of sentences containing them. In fact, the sentences just discussed contained three examples of such words; “I”, “she”, and “now”. Part of learning English involves mastering the context-independent linguistic meaning of these expressions. But mastering the context-independent meanings of these expressions is not sufficient for knowing what they contribute to the truth-conditions of a sentence containing them in a particular context of use. For example, the context-independent meaning of “she” is something like *the indicated (or salient) female*. But in one context, what a use of “she” may contribute to the intuitive truth-conditions of a

sentence such as “She is tired now” is one person, whereas in another context, where another person is being indicated, it may contribute another person. Some knowledge of the extra-linguistic context is needed in order to know the contribution such an expression makes in that context to the intuitive truth-conditions of the sentence containing it. But knowledge of the context-independent meaning of such a term makes the search for such knowledge considerably easier. For example, if a person uses the term “now”, our knowledge of its context-independent meaning restricts our hypotheses about its value to intervals of time containing the moment of that person’s utterance.

My purpose in developing an account of the relation between utterances and their intuitive truth-conditions is to account for two facts. First, we smoothly grasp information about the world from sentences they have never before encountered. Secondly, one and the same sentence may be used to convey different pieces of information about the world, relative to different contexts of use. Both of these facts are unsurprising, given that language is primarily a tool for rapidly conveying information about the world. If we employed a form of communication that lacked context-sensitive elements, then each sentence could only ever be used to express the same piece of information about the world. Such a language, while perhaps helpful for carrying out mathematical investigations, would be impractical for everyday use. It is far easier to utter “that is dangerous”, pointing at a snake in our path, then coming up with a context-insensitive sentence that conveys similar information (perhaps “There is a dangerous snake in front of Jason Stanley and Jeff King”). So for language to be useful, some work has to be done by context-sensitive elements. But if the bulk of intuitive truth-conditions were determined by extra-linguistic context, it would become unclear why we would need to utter sentences at all. My proposal about the relation between utterances and their intuitive truth-conditions is a plausible starting hypothesis for how language is able to be sufficiently elastic as to be usable, and sufficiently rule-governed as to be useful.

There are a number of commitments of the particular metaphysical account I defend of the relation between utterances and their intuitive truth-conditions. It is worth being upfront about them here, rather than forcing the reader to discover them over time in reading through the chapters of this volume.

If every term was context-sensitive in numerous different ways, then, consistently with my proposal, interpretation might still not be systematic. One way this could turn out is as follows. If every term in the language had a very thin context-independent meaning, knowledge of the context-independent meaning would not significantly constrain its interpretation in a context. So, for instance, if the context-independent meaning of “water” was simply *the indicated (or salient) substance*, and the context-independent meaning of “drinks” was simply *the indicated (or salient) relation*, then one would need to know a great deal about the extra-linguistic context in

order to interpret a sentence such as “John drinks water regularly”. For it would be extra-linguistic context that one would need to appeal to discover which substance of all possible substances counted as the content of “water” in that context, and which of all possible relations counted as “drinking” in that context. If so, it might turn out that to grasp the intuitive truth-conditions of an utterance, one had to know a great deal about the extra-linguistic context, far more than is plausible, given the smooth way in which we grasp the truth-conditions of sentences we have not encountered before. One commitment of my project is therefore that, for most words in the language, their meaning is considerably more informative than the above meaning-rules for “water” and “drinks”. This is not a terribly controversial claim, at least not in recent philosophy of language.

The second commitment of my project concerns the effects of extra-linguistic context on the contribution words make, relative to a context, to the truth-conditions of sentences containing them. Even if the context-independent meaning of a word like “water” was fairly rich, and excluded substances such as orange juice and petrol from being water, there is still a worry for the systematic nature of interpretation. If context could affect the interpretation of words in such a manner that the content they express relative to a context could be inconsistent with their context-independent meaning, that would threaten the systematic nature of interpretation. For example, suppose that the context-independent meaning of “water” is the substance pure  $H_2O$ . But suppose that one effect of extra-linguistic context was to allow arbitrarily large expansions of the interpretation of “water”, so that quantities of liquid that are not pure  $H_2O$  could count as water. Then, as far as knowledge of meaning goes, an utterance of “This is water” could be true, relative to some context, just in case the indicated quantity is a quantity of any liquid whatsoever. If context could have this effect on interpretation, then the systematic nature of interpretation would be threatened. As far as meaning goes, any sentence could be associated with virtually any truth-condition in some context.

So another commitment I adopt is that extra-linguistic context is never called upon to *expand* the content determined by the context-independent meaning of a term in a context. For example, if a noun *N* expresses a certain property *P*, then it is never the case that *N* together with the addition made by extra-linguistic context contributes, to the truth-conditions of a sentence containing it, a property that is true of more things than *P*. So, it is plausible to take “tiger” as expressing the property of being a tiger, a certain kind of mammal. A consequence of this second commitment is that there is no context in which the word “tiger”, together with the contributions of extra-linguistic context, contributes, to the intuitive truth-conditions of a sentence containing it, a property such as being a feline, which is true of more things than the property of being a tiger.

The third commitment I adopt involves the significance of syntactic combination. Consider an atomic sentence, such as “Hannah is human”. This sentence predicates “is human” of the content of the name “Hannah”. The semantic significance of predication here may be treated in different theoretical ways. For example, it may be treated as satisfaction (in Tarski’s sense), or functional application (in Montague’s sense), or as property application. So, the truth-condition of “Hannah is human”, relative to a context, is that Hannah satisfies the predicate “is human” or that the function expressed by “is human” maps Hannah onto the true or that Hannah has the property of being human. These different treatments are different ways of explicating the intuitive contribution of predication to the truth-conditions of subject–predicate sentences. But I assume that (in the case of predication) there is one and only one interpretation that is the interpretation of predication in every context. In particular, I assume that the contribution of a syntactic configuration to the intuitive truth-conditions of a sentence containing such a configuration cannot vary as a function of extra-linguistic context, in the way that the contribution of a term such as “she” to the truth-conditions of sentences containing it can vary from context to context.

This last proviso does not exclude the possibility that one syntactic configuration can be ambiguous between two different interpretations, like “bank” is ambiguous between *riverbank* and *financial institution*. What I want to exclude is the possibility that a given syntactic configuration—say the relation between an adjective and a noun in a configuration like “tall woman” or “flat plain”—can have arbitrarily different meanings in different contexts, in the way that a context-sensitive term such as “she” or “I” can have arbitrarily different meanings in different contexts, consistently with its context-independent meaning. In somewhat theoretical terms, this second commitment rules out the possibility of a syntactic configuration having an *indexical* character.

So, this is my proposal for how an utterance comes to be associated with the truth-conditions it intuitively possesses. Each term in a sentence being uttered has a content that is determined by its context-independent meaning together with extra-linguistic context. The context-independent meanings of most terms in the language are fairly rich, rather than fairly vacuous. The function of extra-linguistic context cannot be to expand the content of a term relative to that context. Finally, these contents are put together into truth-conditions by composition rules determined by the syntactic configuration of the sentence, which are not sensitive to context (except in so far as they may be ambiguous between two or three different composition rules). The result is the intuitive truth-condition of the sentence relative to that context (or the intuitive truth-condition of that utterance).

What results from this picture is a very specific metaphysical picture of the way in which utterances become associated with their intuitive truth-conditions. It is trivial to conceive of possible counter-examples. There are many sentences utterances of

which seem to express truth-conditions that cannot be captured in this manner. Let me give some examples.

First of all, consider tense, as in the present-tensed sentence “Bill Clinton is eating lunch”. Intuitively, the truth-conditions of an utterance of this sentence concern the time of utterance. If the utterance of the sentence occurs at 4 p.m. on 12 July 2003, the intuitive truth-conditions of the utterance is that Bill Clinton is eating lunch at 4 p.m. on 12 July 2003. But, unlike the sentence “Bill Clinton is eating lunch now”, there doesn’t appear to be a term in the sentence “Bill Clinton is eating lunch”, the value of which, relative to a context, is the time of utterance. If there is no such term, then utterances of present-tensed sentences show that there are some utterances the intuitive truth-conditions of which are determined by more than just the combination of the contents of the terms in the sentence being uttered.

A second example, as we have seen above, involves predicative uses of gradable adjectives, such as “John is tall” or “Hagia Sophia is old”. Intuitively, an utterance of the sentence “John is tall” may express the content that John is tall for a human being, and an utterance of the sentence “Hagia Sophia is old” may express the content that the venerable Byzantine Church is old for a building. But there does not appear to be a term in either sentence the value of which, relative to a context, is the property of being a human, or (in the second case) the property of being a building. So, this is another example of a kind of sentence, utterances of which appear to have truth-conditions that are not determined just by putting together the contents of the words contained in the sentence in accord with their syntactic configuration.

A third example involves the phenomenon of quantifier domain restriction. Suppose John, having just come back from shopping, utters the sentence “Every bottle is in the fridge”. It seems that the result of adding together the content of each word, relative to this context, in accord with the syntactic structure of the sentence, yields the truth-condition that every bottle *simpliciter* is in the fridge. But, of course, what John intuitively said was that every bottle *that he just bought* is in the fridge. But the restriction to the things John just bought does not seem to come from the contents of any part of the sentence uttered. So, this is another case where my proposal as to the relation between a sentence and the intuitive truth-conditions of an utterance of it seems incorrect.

A fourth example involves mass-terms, such as “water” or “gold”. It is plausible to take a term like “water” to express the property of being pure H<sub>2</sub>O. If so, then the contribution of “water” to the intuitive truth-conditions of utterances of sentences containing it must be the property of being pure H<sub>2</sub>O. But consider an utterance of “Lake Burley Griffin is filled with water”. Intuitively, this utterance is true. So, its intuitive truth-conditions must in fact be satisfied. But given the above hypothesis about the content of “water”, the truth-conditions of this utterance are not actually met. For it is not true that Lake Burley Griffin is filled with pure H<sub>2</sub>O. The substance

that fills Lake Burley Griffin is not pure H<sub>2</sub>O, but has H<sub>2</sub>O mixed with many impurities. So this is another example that appears to threaten a part of my project. For here it seems that the contribution of “water” to the intuitive truth-conditions of utterances of sentences containing it is an “expansion” of its stand-alone content. That is, the contribution of “water” to the intuitive truth-conditions of utterances of sentences containing it is a property that is true not just of quantities of pure H<sub>2</sub>O, but also of H<sub>2</sub>O mixed with impurities.

There are numerous other examples that suggest that my proposal about the relation between an utterance and its intuitive truth-conditions is overly restrictive. Some may think of this as a reason not to pursue it at all. But my own inclination is to seek an account of the relation between utterances and their intuitive truth-conditions that makes the connection between them explicable. I have great difficulty seeing alternative proposals as sufficiently constrained. Proposals that appear smoothly to account for the bewildering variety of ways in which context appears to affect the relation between utterances and their intuitive truth-conditions only do so because they appeal (whether overtly or covertly) to interpretive processes that are deeply mysterious or wildly unconstrained. I prefer to explore a view that is neither mysterious nor unconstrained, and see how far one can use it to explain our linguistic behavior. If, at the end of the day, my proposal is too constrained, perhaps the investigation into its adequacy will reveal systematic ways of modifying or liberalizing it. This seems more promising than beginning with a mysterious and unconstrained process, and trying to add on stipulative constraints. If we do not have a clear grip on the process with which we began, why think that adding on restrictions will yield any greater elucidation?

Fortunately, the alternatives are not “explicable but subject to counter-example” versus “inexplicable, but not subject to counter-example”. The proposal I have sketched above concerning the relation between utterances and their intuitive truth-conditions may seem at first blush to face insurmountable objections. But I argue in this book that a host of prima-facie counter-examples to the proposal are not successful. For example, I argue, in each of the above cases, that the intuitive truth-conditions of the utterance are due to the assignment of values to the parts of the sentence uttered, and combination in accord with syntactic structure. In each case, the belief that this cannot be the case is due either to an impoverished conception of syntactic structure, or an impoverished conception of available semantic resources.

Many philosophers have been driven to reject the kind of intimate relation I favor between the result of combining the contents of parts of a sentence in a context and the intuitive truth-conditions of its utterance because of their antecedent commitments about certain controversial constructions. For example, consider true identity sentences flanked by names, e.g., “Hesperus is Phosphorus” or “Cicero is Tully”. Some philosophers of language hold that the intuitive truth-conditions of



propositional attitude ascriptions containing such sentences (such as “John believes that Cicero is Tully”, or “Hannah doubts that Hesperus is Phosphorus”) diverge quite radically from the result of taking the contents of their parts and combining them in accord with their syntactic structure. For such philosophers, the result of the latter procedure, applied to a sentence like “John believes that Cicero is Tully”, yields the same content as the result of the procedure, applied to the sentence “John believes that Cicero is Cicero”. But, intuitively, it certainly seems that an utterance of “John believes that Cicero is Tully” could be false, even though John believes that Cicero is Cicero.

I do not, in the chapters in this volume, address propositional attitude constructions, and in general (except for the case of quantifier domain restriction and indicative conditionals) I steer clear of constructions that have engendered the greatest philosophical controversy. But that does not mean the proposal is not relevant for these topics. If the proposal works outside the realm of the most philosophically contentious cases, then that provides strong evidence that it is correct in the philosophically contentious cases. Thus the success of this project would provide motivation for investigation of views of propositional attitude constructions of the sort defended by Peter Ludlow, according to which there are unpronounced variables in the scope of propositional attitude constructions whose values are modes of presentation, ways of thinking, or other non-extensional entities.

The first chapter, “Context and Logical Form”, introduces my proposal, and argues, in a preliminary fashion, that a wide variety of constructions can be accommodated within its scope. The first section is devoted to explaining the proposal and its commitments. In the second section, I consider the apparent phenomenon of “non-sentential assertion”. Those who reject my account of the relation between utterances and their intuitive truth-conditions often appeal to the fact that we make assertions (thereby presumably expressing truth-conditions) with the use of apparently sub-sentential items. For example, a thirsty man might utter “water” to a street vendor, or I might utter “two apples” to someone selling apples. If such utterances are genuinely non-sentential, and genuinely are assertions of truth-conditions, then we have strong evidence that the relation between an utterance and its intuitive truth-conditions can be mediated by processes quite unlike assignments to parts of a sentence and combination. If this phenomenon is genuine, then something like an existence proof has been provided for “free” contextual enrichment (cf. Carston (2002, 63)) of the sort discussed above. In this second section, I argue that the phenomenon is not genuine or, more precisely, that to describe it as “non-sentential assertion” is incorrect. Finally, in the last section, I consider a wide variety of apparent counter-examples to my proposal about the relation between utterances and their intuitive truth-conditions, and argue that in fact they involve sentences whose structure is more complex than it appears to be on the surface.

The purpose of “Context and Logical Form” was to introduce the general project, and defend some of the methodology used to establish hidden structure. Discussion of particular examples is therefore not pursued in great detail. The purpose of the second chapter “On Quantifier Domain Restriction”, co-authored with Zoltán Gendler Szabó, was to investigate one case in great detail, that of quantifier domain restriction. When someone utters a quantified sentence, such as “Every bottle is in the fridge”, they usually intend to communicate something about a restricted domain of things. In this chapter Szabó and I argue that one can explain the relation between an utterance of a quantified sentence and its intuitive truth-conditions within the restrictions imposed by my proposal in “Context and Logical Form”. In particular, we argue that quantifier domain restriction is due to the presence of unpronounced structure in the structure of the quantified sentence uttered.

The hidden structure postulated in “On Quantifier Domain Restriction” was not entirely motivated by the arguments provided in that paper. For example, in that paper, following work by Kai von Stechow (1994), we argued that a domain variable is really a package of a function variable and an object variable.<sup>2</sup> The object variable is assigned objects, relative to a context, and the function variable is assigned a function from objects to properties (we also argue that domains are properly conceived of as properties rather than sets). So, in a sentence such as “Every fireman is tough”, there are really two domain variables, which together determine the quantifier domain for the quantified noun phrase “every fireman”. Relative to a context, the variable ranging over objects associated with “every fireman” is perhaps assigned a location, and the variable over functions is perhaps assigned a function from locations to the property of being an inhabitant in that location. Examples by von Stechow and Robin Cooper (such as “In every township, every fireman is tough”) show that the object variable can be bound by a higher quantifier, which provides evidence of a syntactic sort for the existence of that variable. But no one had provided similar evidence of a similar sort for the existence of the function variable that is involved in domain restriction. For instance, no one had provided examples in which the function variable associated with domain restriction can be bound by a higher quantifier, despite the fact that the function variable is required for the semantic analysis to work correctly. In “Nominal Restriction”, the third chapter in this volume, I provide the missing motivation for the hidden structure postulated by Szabó and my analysis. I then argue that the existence of this structure explains not only the phenomenon of quantifier domain restriction but also a number of other apparently distinct contextual effects.

One kind of apparently distinct contextual effect explained by the apparatus of quantifier domain restriction involves the context-sensitivity of adjectival constructions. As Hans Kamp pointed out many years ago, the sentence “John built a large

<sup>2</sup> Von Stechow associates these variables with the determiner rather than (as we do) with the noun.

snowman” is intuitively context-sensitive. Suppose that John has been injured recently and, as part of his therapy, has been building snowmen. On this particular occasion, he has built a snowman that is much larger than his previous snowmen. Then an utterance of “John built a large snowman” may express a truth, even though the snowman he built is not particularly large for a snowman built by a healthy person of his age. As it turns out, recognizing that nouns are associated with domain indices allows for a solution to this problem.

Much of “Nominal Restriction” concerns context-sensitivity associated with mass terms. As with count nouns such as “bottle” and “sailor”, mass nouns can occur in quantificational structures. So, when I declare that there is a little milk in the refrigerator, my utterance is not made true by the existence of a small drop of milk in the corner; intuitively “little milk” ranges only over larger quantities of milk. The nominal restriction theory advocated by Szabó and me explains this phenomena; “milk” is associated with a domain index, restricting the quantification to quantities of a certain size. But it also explains other kinds of context-sensitivity associated with mass terms, context-sensitivity more typically explained by appeal to “loose use”. Many philosophers think that “water” denotes quantities of pure H<sub>2</sub>O. When we say that lakes and rivers contain water, we are speaking loosely (since the stuff that flows in lakes and rivers is filled with impurities). But note that count nouns do not allow a similar kind of loose use; we never describe housecats as impure tigers. This suggests to me that the idea that mass terms such as “water” and “gold” have, as their extensions, only quantities of pure chemical kinds, is mistaken. Rather, what we have here is a species of semantic context-sensitivity, rather than loose talk. In the final section of “Nominal Restriction”, I argue that the nominal restriction theory justifies this intuition; terms such as “water” and “gold” contain, in their extensions, quantities of impure substances. In contexts in which we speak of only H<sub>2</sub>O being water, we are employing domain restriction.

Many debates in the philosophy of language are confused by the fact that theorists employ distinct conceptions of the relation between semantics and pragmatics. In “Context and Logical Form”, I was very specific about my use of term “semantics”, and I distinguished it from the use of the term “semantics” employed by some other theorists, such as John Perry. But, under the urging of Jeffrey King, I soon came to realize that a more thorough discussion and justification of terminological choices was necessary. Out of our discussions of these matters came our co-authored chapter “Semantics, Pragmatics, and the Role of Semantic Content”, much of which is devoted to an overview of different conceptions of the semantics–pragmatics distinction.

The first part of the chapter is devoted to justifying our favored way of drawing the semantics–pragmatics distinction. For example, some theorists choose to extrude from the semantic content of sentences relative to contexts any effects of speaker intentions. According to such theorists (most prominently Kent Bach, but also

Emma Borg), the semantic content of “That is a table” is not a proposition, because the reference of “that” is determined by speaker intentions, and speaker intentions do not influence semantic content, even relative to a context. Presumably, one motivation for extruding the effects of speaker intentions from semantic content is the aforementioned concern that adding context-sensitivity to the semantics would make the semantics unsystematic. We argue in this chapter that there are certain clear ways in which not recognizing the semantic effects of extra-linguistic context (including perhaps speaker intentions) could make the semantics less systematic, by posing a problem for compositionality.

The final part of our joint chapter is another defense of the central thesis of the essays in this book, namely that the scope of semantic content includes intuitive truth-conditions. King and I address a host of constructions that classically have been held to pose a threat to those who wish to tie intuitive truth-conditions closely to sentential content (conceived of as on my proposal). For example, consider the comparative “Eating some of the cake is better than eating all of it”. Intuitively, an utterance of this expresses the truth-condition that eating some but not all of the cake is better than eating all of the cake. But the result of composing the contents of the parts of the sentence “Eating some of the cake is better than eating all of it” does not appear to yield this intuitive truth-condition. Certainly, the content of “some of the cake” is not “some but not all of the cake”—one way to eat some of the cake is to eat all of it. Similarly, consider a conditional such as “If John marries Hannah and has kids, he will be happy, but if he has kids and marries Hannah, he won’t”. Intuitively, an utterance of this conditional has the intuitive truth-condition that if John marries Hannah and subsequently has kids, he will be happy, but if he has kids and subsequently marries Hannah, he won’t. But this content does not seem derivable from the contents of the words occurring in the conditional construction. In this chapter we argue that none of these classic problem cases poses worries for my proposal about the relation between utterances and their intuitive truth-conditions. The claim that these cases do raise worries for my proposal derives from an overly simplistic view of the semantics of terms such as “if . . . then” and “better than”.

The next chapter in this volume is “Making it Articulated”. The central purpose of this chapter is to show that the non-semantic explanations others have suggested to accommodate data about speakers’ intuitions about the truth-conditions of their utterances are uniformly unsuccessful. The worries with my proposed account of the relation between sentences in context and their intuitive truth-conditions all involve under-generation. That is, my proposal about the relation between utterances and their intuitive truth-conditions is highly restrictive, and there are many examples of utterances that seem to have intuitive truth-conditions that cannot be captured by my proposal. In contrast, the processes that are appealed to in the pragmatics literature have the opposite problem; they tend to *over-generate*, predicting the availability

of readings of certain utterances that simply are not there. My purpose in this chapter is to emphasize the depth of the over-generation problem facing attempts to accommodate the data in other ways.

The way I bring out the over-generation problem facing pragmatic explanations of apparently semantic phenomena is by focusing on attempted pragmatic explanations of data that I account for by appeal to hidden syntactic structure. In particular, my focus is on “free enrichment” accounts of *bound* readings of certain sentences, such as “Every species has members that are old”, where the comparison class for “old” seems to be bound by the initial quantifier “every species”. In “Context and Logical Form”, I had argued that this example provides evidence for unpronounced structure associated with predicative adjective constructions. I had also (following others, such as Kai von Stechow) argued for a similar conclusion for the case of quantifier domains. Advocates of free pragmatic enrichment had responded to my arguments by appeal to free pragmatic enrichment with variables. My purpose in “Making it Articulated” is to bring out the severity of the over-generation problem facing advocates of pragmatic accounts of this kind of data, by criticizing free pragmatic enrichment with variables. As I argue in this chapter, those who appeal to free pragmatic enrichment to explain away bound readings of sentences are at a loss to explain why many sentences lack readings they should have, if free pragmatic enrichment were an interpretive option.

Note that on my own proposal, what is said by an utterance of a sentence does require a good deal of what is traditionally called “pragmatics”. For, on my view, each sentence contains a number of unpronounced context-sensitive constituents. For example, understanding what is said by an utterance of “Every bottle is in the fridge” requires grasping the value of domain indices associated with “bottle”. This involves grasping the contents of both a function variable and an object variable. Just as one cannot “read off” the proposition expressed by an utterance of “That is green” from the sentence used, one cannot “read off” the proposition expressed by an utterance of “Every bottle is in the fridge” from the sentence used. The sentence used restricts the possibilities of interpretation in very specific ways, thereby easing the interpretive problem facing the hearer. But the multiplicity of context-sensitive elements, both pronounced and unpronounced, nevertheless places demands on both speaker and hearer. The hearer must infer the values of the context-sensitive elements in the utterance she apprehends, or something close enough to them, and the cooperative speaker must ensure that her audience is in a position to grasp them, or something close enough to them.

So the difference between my proposal and proposals that postulate an interpretive process of free enrichment is *not* that the latter kinds of approaches require participants in a conversation to engage in a great deal of reasoning about the context to grasp what is intuitively expressed. My proposal also requires interlocutors to engage in reasoning about the context. Such reasoning is needed in order to resolve the values

of context-sensitive elements in the sentence, and to make sure that one's hearers can easily resolve them. My proposal embodies a quite specific hypothesis about *how* sentences constrain the possibilities of their correct interpretations, one that is considerably more restrictive than proposals that postulate "free enrichment", yet promises to explain all the data without facing the threat of over-generating readings of sentences.

Chapter 6 of this volume is "Semantics in Context". I begin by setting out the challenge from context-sensitivity, and discussing some of the most difficult problem cases for a view such as my own, which seeks to reflect the intuitive truth-conditions of an utterance in the semantic content for the sentence uttered, relative to a context of use. One such difficult problem case is the problem of *deferred reference*. Consider, for example, an annoyed waitress muttering "The ham sandwich is irritating", or a doctor telling a nurse, "The kidney in room 103 needs a glucose drip". The intuitive truth-conditions of the first utterance are that the person who ordered the ham sandwich is irritating, and the intuitive truth-conditions of the second utterance are that the person with the kidney problem in room 103 needs a glucose drip. The problem with the phenomenon of deferred reference is that it does seem to affect the intuitive truth-conditions of an utterance, and yet incorporating it into the theory of meaning proper threatens the systematic nature of semantic theory (or so I maintain). In sections I and II of the chapter, I try to develop several principled reasons to think that the phenomenon of deferred reference is not semantic in nature. In so doing, I try to give criteria for when a phenomenon that seems to affect intuitive truth-conditions should be incorporated into a semantic account of what is said, criteria that are consistent with the demand that a semantic theory must be responsive to facts about speaker intuitions about truth-conditions of utterances.

In the third, fourth, and fifth sections of the chapter, I return to a discussion of one piece of methodology that I have used to establish the presence of hidden structure, namely the binding argument, which involves an inference from the existence of bound readings of a sentence to the existence of unpronounced variables in the logical form of the relevant construction. So from the fact that "Every species has members that are old", I conclude that a predicative adjective construction such as "is old" involves an unpronounced variable associated with "old". One way to reject this inference is to appeal to free pragmatic enrichment; the bound reading is explained by the fact that speakers have available to them an interpretive strategy that can add variables to sentences they hear. In "Making it Articulated", I discuss and reject such accounts. My purpose in sections IV and V of "Semantics in Context" is to evaluate two alternative accounts that also would block the inference from the existence of bound readings to the existence of unpronounced variables in logical form. Unlike free pragmatic enrichment, these alternative accounts are each semantic in nature.

But I argue that neither suffices to account for the presence of bound readings in all cases.

In “Making it Articulated”, I level the charge of over-generation against advocates of free pragmatic enrichment. What I mean by “over-generation” in that chapter is that such accounts over-generate *readings of sentences*; if free pragmatic enrichment were an interpretive option, there should be readings of sentences that simply are not interpretive options for utterances of those sentences. But there is a somewhat related over-generation concern facing my own proposal. It is not that my account predicts there to be certain non-existent readings of sentences; it is rather that, given the interpretive possibilities, my account requires too much hidden structure. The final section of “Semantics in Context” is devoted to countering this concern.

Since I began this project many years ago, a number of authors have published books developing their views. One of the most important contributions to this literature is François Recanatì’s *Literal Meaning*. I include my review of this book, in order to allow the reader some sense of the relation between my views and opposing positions.

## References

- ALMOG, JOSEPH, PERRY, JOHN, and WETTSTEIN, HOWARD (eds.) (1989) *Themes from Kaplan* (Oxford: Oxford University Press).
- BACH, K. (1994), ‘Conversational Implicature’, *Mind and Language*, 9, 124–62.
- BEZUIDENHOUT, ANNE (1997), ‘The Communication of De Re Thoughts’, *Nous* 31/2: 197–225.
- BORG, EMMA (2004), *Minimal Semantics* (Oxford: Oxford University Press).
- CAPPELEN, HERMAN, and LEPORE, ERNIE (2004), *Insensitive Semantics: In Defense of Semantic Minimalism and Speech Act Pluralism* (Blackwell).
- CAPPELEN, HERMAN, and LEPORE, ERNIE (2006), ‘Shared Content’, in E. Lepore and B. Smith (eds.), *Oxford Handbook of Philosophy of Language* (Oxford University Press).
- CARSTON, ROBYN (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication* (Oxford: Blackwell).
- KAI VON FINTEL (1994), *Restrictions on Quantifier Domains*, Ph.D. thesis (Amherst: University of Massachusetts).
- GIBBARD, ALLAN (1981), ‘Two Recent Theories of Conditionals’, in W. L. Harper, R. Stalnaker, and G. Pearce (eds.), *Ifs* (Dordrecht: Reidel) 211–47.
- GRICE, PAUL (1989), *Studies in the Way of Words* (Cambridge: Harvard University Press).
- GRICE, PAUL (1989a), ‘Logic and Conversation’, in Grice (1989), 22–40.
- HAWTHORNE, JOHN (2004), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- HECK, RICHARD (2002), ‘Do Demonstratives Have Senses?’, *Philosophers’ Imprint*, 2.
- KAPLAN, DAVID (1989), ‘Demonstratives’, in Almog, Perry, and Wettstein (eds.) (1989), 481–565.
- LUDLOW, PETER (1995), ‘Logic Form and the Hidden Indexical Theory: A Reply to Schiffer’, *Journal of Philosophy*, 92: 102–7.

- SPERBER, D., and WILSON, D. (1986), *Relevance* (Cambridge: Harvard University Press).
- STALNAKER, ROBERT (1987), *Inquiry* (Cambridge: MIT Press).
- STRAWSON, PETER (1996), 'On Referring', in A. P. Martinich (ed.), *The Philosophy of Language*, (New York: Oxford University Press, 3rd edn.), 208–30.
- TIMOTHY WILLIAMSON (2005), 'Knowledge, Context, and the Agent's Point of View', in Preyer and Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 91–114.



# 1

## Context and Logical Form

My purpose in this chapter is to defend the thesis that all truth-conditional effects of extra-linguistic context can be traced to logical form.<sup>1</sup> But before the import of this thesis can be understood, a few distinctions must be clarified, and some of its opponents introduced.

There is certainly no one uniform use of the expression “Logical Form”. But there are two distinguishable senses underlying its many differing usages. It is only in the second of these two senses that the thesis I will defend is interesting and controversial.

Perhaps the most prevalent tradition of usage of the expression “Logical Form” in philosophy is to express what one might call the *revisionary* conception of Logical Form. According to the revisionary conception, natural language is defective in some fundamental way. Appeals to Logical Form are appeals to a kind of linguistic representation which is intended to replace natural language for the purposes of scientific or mathematical investigation. Different purposes may then give rise to different regimentations of natural language. For example, one might want to replace natural language by a notation in which there is some kind of isomorphism between the true sentences in the notation and the facts they describe (e.g., Russell (1985)). Alternatively, one might want to replace natural language by a notation which explicitly reveals the hidden contribution of logical expressions, such as the language of the predicate calculus.

To say that all context-dependence is traceable to logical form in a revisionary sense of “logical form” might be taken to be the trivial claim that, for purposes of interpretation, one should replace natural language by a notation in which all context-dependence is made explicit in the favored notation. It is not in this sense that I intend the thesis.

According to the second tradition of usage, which one might call the *descriptive* conception of Logical Form, the Logical Form of a sentence is something like the “real structure” of that sentence (e.g., Harman (1972)). On this approach, we may

<sup>1</sup> By “context” in this chapter I will throughout mean extra-linguistic context. So, nothing I say bears on standard appeals to type-shifting principles, which involve the effects of linguistic context on interpretation.

discover that the “real” structure of a natural language sentence is in fact quite distinct from its surface grammatical form. Talk of Logical Form in this sense involves attributing hidden complexity to sentences of natural language, complexity which is ultimately revealed by empirical inquiry. It is in this sense that I intend the thesis that all context-dependence is traceable to logical form. What I will defend is the claim that all truth-conditional context-dependence results from fixing the values of contextually sensitive elements in the real structure of natural language sentences.

## I

In this chapter I focus, for clarity’s sake, on the speech act of assertion. My goal will be to defend the claim that all effects of extra-linguistic context on the truth-conditions of assertions are traceable to logical form. Though ordinary language philosophers (e.g., Austin (1970)) held such generalizations to be illegitimate, I will nevertheless assume here that the arguments I advance for the case of assertion generalize to other speech acts. I will also assume that each successful assertion has a truth-condition. I will often call the truth-conditions of an assertion “what is expressed by that assertion”. This usage must be sharply distinguished from the usage found in authors such as Bach (1994), where “what is expressed” is allowed to denote something that is not a truth-condition.

At times in the course of this chapter I will speak of a truth-condition as a certain kind of thing, namely a *structured proposition*, an ordered sequence of objects and properties.<sup>2</sup> There are two reasons for this. First, many philosophers think of a semantic theory for a language as primarily involving an algorithm which assigns structured propositions to sentences relative to contexts, and so are more familiar with the issues when couched in these terms. Perhaps also because of the first reason, appeal to talk of structured propositions makes the issues I discuss somewhat simpler to explain. For example, it allows us to speak of “constituents” of what is expressed that correspond to constituents of sentences. It is slightly more difficult to avail ourselves of this useful metaphor on a straight truth-theoretic framework.<sup>3</sup> For these

<sup>2</sup> In his (1987), Scott Soames contrasts his structured proposition conception of semantics with a truth-conditional conception, and rejects the latter on the grounds that truth-conditions are too fine-grained to serve fundamental semantic purposes. However, these (important) issues are independent of the concerns of this chapter. If one thinks of structured propositions as more fundamental entities that determine the truth-conditions of an assertion, then one can take this chapter to concern the proposition expressed by an assertion, and only derivatively the truth-conditions of an assertion.

<sup>3</sup> One cannot, in a truth-theoretic semantics, speak of the “constituents” of what is expressed that correspond to the sentence. Rather, one must speak, more awkwardly, of the properties and objects

two reasons, I will occasionally speak in these terms, though nothing substantial rests upon my uses of this framework.

I will also assume in this chapter that syntax associates with each occurrence of a natural language expression a lexically and perhaps also structurally disambiguated structure which differs from its apparent structure, and is the primary object of semantic interpretation. In accord with standard usage in syntax, I call such structures *logical forms*.

I will repeatedly be using the terms “semantic” and “pragmatic”. However, there are many different usages of these expressions. For example, according to one traditional use of the term “semantic”, semantics is the study of context-invariant aspects of meaning. On this account, the semantic content of any two utterances of “I am tired” is the same, since their context-invariant meaning is the same. If one is using the term “semantic” in this sense, then there is a corresponding sense of the term “pragmatic”. In this sense, pragmatics is the study of those aspects of linguistic communication that depend on context. For example, the study of how the meaning of indexical expressions changes with their context of use is, on this way of using the terms, part of pragmatics (cf. Bar-Hillel (1954)). Before the work of Paul Grice, this use of “semantic” and “pragmatics” was standard. For example, it seems to be the best explication of the usages of the terms in the work of Richard Montague.

This usage of the expressions “semantic” and “pragmatics” is very clear. However, it obscures important disanalogies. It is very natural to divide the process of linguistic interpretation into two phases. In the first phase, a hearer first assigns denotations to each element of the logical form produced by the speaker, denotations that are determined by the meanings of the elements of the logical forms plus perhaps contextual factors. The hearer then combines these values in accordance with the structure of the logical form to derive the interpretation of the logical form, relative to that context. In many cases, for example, words like “I”, “now”, “this”, and “she”, the context-invariant meaning of an element in the logical form does not exhaust its denotation, but rather serves a guide for the interpreter in this process of denotation assignment. In the second phase, the hearer evaluates the result of the first phase with respect to general conversational maxims, such as relevance, quality, or quantity. This second stage of interpretation is not linguistic in nature. It does not involve the assignment of values to elements of a structured representation produced by the speaker. Accordingly, the first stage of interpretation is “semantic”, the second, “pragmatic”.<sup>4</sup>

These two usages are very different. According to the first usage, what semantics interprets are expression types, *simpliciter*. On the first usage, there are no semantic

mentioned in the statement of the truth-conditions that are introduced by rules assigning them to expressions of the object-language.

<sup>4</sup> Bach (ms) gives a particularly clear explanation of this notion of “semantic”.

differences between distinct uses of a sentence such as “I am tired”. According to the second usage, on the other hand, what semantics interprets are rather *expressions relative to contexts*. If Hannah is the speaker in context *c*, and John is the speaker in context *c'*, then there is a semantic difference between “I am tired”, relative to *c*, and “I am tired”, relative to *c'*.

There is a third very standard usage of “semantic” and “pragmatic”. According to this third usage, semantics concerns truth-conditions, or propositions. There are many different usages of this familiar phrase (cf. Stalnaker (1970), for one such usage). However, the usage I have in mind is one according to which the phase of interpretation that is semantic is the one that results in truth-conditions (see pp. 79–80). It is this usage that underlies talk of “truth-conditional semantics”. Pragmatics is then the study of those aspects of interpretation that take as input the truth-conditions of a linguistic act, and yield other propositions implicated by that speech act. This is the usage that is most clearly suggested by the work of Grice.<sup>5</sup>

These three distinct usages do not come close to exhausting the different senses of “semantic” and “pragmatic” in the literature. To avoid debates that are at bottom terminological, it is important, in any discussion of issues involving context, to settle immediately upon one way of using these expressions. In this chapter I use the expressions “semantic” and “pragmatic” in the second of the above senses. That is, semantic interpretation involves the assignment of denotations to elements of a logical form relative to a context, and their combination. Extra-linguistic context enters in only when called upon by a linguistic rule governing an element. The result of semantic interpretation is some kind of non-linguistic entity, such as a proposition or a property, which is then the input to pragmatics.

However, if my claim in this paper is correct, then the second and third usages of “semantic” and “pragmatic” coincide. That is, if all effects of extra-linguistic context on the truth-conditions of an assertion are traceable to logical form, then the result of semantic interpretation in the second of the above usages will be the truth-conditions of the assertion, and hence the result of semantic interpretation, in the third sense of “semantic”. One purpose of my chapter is an attempt to bind together these two distinct usages of the term “semantic”, and thereby justify talk of “truth-conditional semantics”.

<sup>5</sup> For example, Grice is very clear that his “favored use” of “what is said” applies to utterances, or expressions in contexts, rather than expression types. Considering an utterance of “He is in the grip of a vice”, made about some person *x*, Grice writes “for a full identification of what the speaker said, one would need to know (a) the identity of *x*, (b) the time of utterance, and (c) the meaning, on the particular occasion of utterance, of the phrase ‘in the grip of a vice’” (1989, 25). As this passage makes clear, there is also no reason to think that Grice thought that every element of what is said must be the value of something in the logical form, since he claims that the time of utterance is a determinant of what is said, but never suggests that it is named by a constituent of the sentence. It is what is said in Grice’s favored sense that is, according to him, the input to pragmatics.

Now that we are clear about my future use of the term “semantic”, I add a final assumption. The assumption is that composition rules do not vary as a function of extra-linguistic context. This assumption is entailed by every version of the principle of compositionality, which is a standard condition of adequacy on a semantic theory. According to one formulation of this principle, a semantic theory is compositional just in case, for each complex expression, there is exactly one way, determined solely by its structure, in which the meanings of its constituents are combined by the semantic theory to yield its meaning.<sup>6</sup> It follows from this principle that, although the meaning of a non-complex word may vary with context, the way in which the interpretation of a complex expression is derived from the interpretations of its parts cannot vary with context. For if a semantic theory allowed the way in which the interpretation of a complex expression is built from the interpretation of its parts to vary with context, then it would not correlate with each complex expression, a *unique* way in which the interpretation of its constituents combine to yield its interpretation.<sup>7</sup>

There are certain authors who reject the principle of compositionality, since they hold that the meaning of a complex expression may depend upon its *linguistic* context (e.g., Higginbotham (1986), Hintikka and Sandu (1997)). However, this position is fully consistent with the assumption I have made, that composition rules do not vary as a function of *extra-linguistic* context. This latter assumption is far weaker than compositionality. A semantic theory which violates it would, I suspect, simply be unlearnable. Since I am not aware of any author who is not an opponent of systematic semantics who has denied it, the assumption should be uncontroversial, and I presuppose it in what follows.

Suppose my principal claim is true, that all effects of extra-linguistic context on the truth-conditions of an assertion are traceable to logical form. Then, the effects of context on the truth-conditional interpretation of an assertion are restricted to assigning the values to elements in the expression uttered. Each such element brings with it rules governing what context can and cannot assign to it, of varying degrees

<sup>6</sup> On this characterization of compositionality, different syntactic constructions may be associated with different modes of semantic composition. For a useful discussion of different notions of compositionality and related principles, see Janssen (1997).

<sup>7</sup> It is worth mentioning that most semantic accounts of variable-binding are in tension with compositionality as I have stated it (though of course are consistent with my assumption). An example is the “Predicate Abstraction Rule” discussed in Heim and Kratzer (1998, 186 ff.). Essentially, Heim and Kratzer assume a syntax that involves structures such as:  $[\alpha \lambda x [S [NP [N \text{John}]] [VP [V \text{offended}] [NP [N x]]]]]$ . However, they assign no independent interpretation to “ $\lambda x$ ”. So, on their account, interpreting the node  $\alpha$  does not amount to combining the value of “ $\lambda x$ ” with the semantic value of the open sentence “John offended  $x$ ”. Rather, they provide a non-compositional interpretation rule. This sort of violation of compositionality is fairly common, and should not raise any worries. Violations of compositionality only become worrisome from the standpoint of learnability when they involve an unlimited number of unrelated construction rules, as would be the case with context-dependent construction rules.

of laxity. The effects of extra-linguistic context on truth-conditional interpretation are therefore highly constrained. If this picture of truth-conditional interpretation is correct, then it is fundamentally different from other kinds of interpretation, like the kind involved in interpreting kicks under the table and taps on the shoulder.<sup>8</sup> We do not interpret these latter sorts of acts by applying highly specific rules to structured representations. Nor is the role of extra-linguistic context in interpreting these acts in any way constrained, as it is in the case of linguistic interpretation. Thus, if the interpretation of assertions in fact functions in the way I have sketched, one should be suspicious of views that assimilate it too quickly to the ways in which we interpret non-linguistic acts.

In recent years, there has been no shortage of philosophers of language and cognitive scientists eager to reject the claim I have advanced. According to Kent Bach, Robyn Carston, François Recanati, Dan Sperber, Robert Stainton, Charles Travis, and Deirdre Wilson, among others, the truth-conditions of most assertions go well beyond what semantics can legitimately assign to the logical forms of the sentences uttered. Instead of assigning propositions, entities that are truth-evaluable, to logical forms, semantic interpretation only involves “fragmentary representations of thought” (Sperber and Wilson (1986, 193)) “partially articulated conceptual representations” (Carston (1991, 49)), or “propositional radicals” (Bach (1994, 127); cf. also Bach (1982)).<sup>9</sup> The examples motivating these theorists all concern the effects of context on what is expressed in assertions. According to these theorists, there is no way to “constrain” the effects of context on what is expressed within the domain of semantic interpretation. In most cases, what the semantic interpretation of a sentence’s logical form delivers is not what is expressed, but rather, in the words of Sperber and Wilson, “mental objects that never surface to consciousness”; these are then used in a *pragmatic* derivation of what is expressed.

If these theorists are correct, then semantics is not about truth-conditions. It would then be more apt to replace, as does Recanati, talk of truth-conditional *semantics* with talk of truth-conditional *pragmatics* (cf. Recanati (1993, ch. 13)).

Underlying these arguments against the picture of interpretation I advocate are two assumptions about semantic theory, both of which I accept. The first assumption these theorists make about what is semantically legitimate is:

**First assumption:** In semantic interpretation, one may never postulate hidden structure that is inconsistent with correct syntactic theory.

<sup>8</sup> Where the latter are not governed by explicit meaning-granting stipulations. This proviso should be tacitly understood in future references to interpretation of non-linguistic acts.

<sup>9</sup> Note that the first two ways of speaking involve something like use/mention errors according to the usage of semantics at issue in this chapter. To talk of semantic interpretation resulting in “representations” seems *prima-facie* confused, for the reasons discussed in section I of Lewis (1983).

According to some conceptions of semantics, the objects of semantic interpretation are not syntactic logical forms, but rather logical forms in some more revisionary sense of “logical form”. With a revisionary conception of logical form, one is not constrained by the actual syntactic structure of the sentences under consideration. A theorist in this tradition could reject the first assumption, noting that her postulated hidden structures are not intended to be indicative of the actual syntactic structures of the sentences used. However, together with the advocates of truth-conditional pragmatics, I reject this conception of semantics. The objects of semantic interpretation are the actual logical forms of English sentences; the first assumption is simply a consequence of this.

The second assumption appears under various names in the literature, such as the “linguistic direction principle” (Carston (1991, 38–9)) or the “criterion of close syntactic correlation” (Bach (1994, 137)):

**Second assumption:** In deriving the semantic interpretation of a logical form, every feature of the semantic interpretation must be the semantic value of something in that logical form, or introduced via a context-independent construction rule.

This assumption is also clearly part of the conception of semantics I have articulated above.

Here is how the proponents of truth-conditional pragmatics use the two assumptions to argue that truth-conditions are not determined by semantics. First, some linguistic construction is provided whose truth-conditional interpretation is mediated by context. Then, it is argued that it is inconsistent with current syntactic theory to postulate, in the logical form of the relevant construction, expressions or variables the semantic values of which context could provide. So, by the first assumption, it follows that the information provided by context to the truth-conditional interpretation of the relevant construction is not the semantic value of anything in the syntactic logical form. By the second assumption, it then follows that the information provided by context to the truth-conditional interpretation of the construction is not a part of semantic interpretation at all. The conclusion is that, in such cases, semantic interpretation does not deliver truth-conditions.

Given these two assumptions, it is an empirical question whether there are constructions whose truth-conditional interpretation is not entirely a matter of semantics. However, it is an empirical question whose resolution has significant foundational consequences. If the advocates of truth-conditional pragmatics are correct, then the proper place to situate an account of the bulk of the truth-conditional interpretation of linguistic assertions is in whatever account one has of reasoning generally, regardless of its subject-matter. If, by contrast, the truth-conditional interpretation of assertions is entirely a matter of semantics, then the

truth-conditional interpretation of assertions is special in a way that other kinds of reasoning processes are not.<sup>10</sup>

There are essentially two lines of response available to the opponent of truth-conditional pragmatics. The first is to reject the conception of semantics I have adopted, allowing that semantic interpretation is not just interpretation of the words used. There are no doubt different ways to accomplish this. For example, according to J. L. Austin, it is statements, *acts* of asserting sentences, of which truth is ultimately predicated (cf. Austin (1979)). According to this alternative conception of semantics, championed by theorists such as Mark Crimmins and John Perry, it is utterances that are the ultimate objects of interpretation, rather than the sentences uttered in these acts. On this approach, since it is not logical forms relative to contexts that are the ultimate objects of interpretation, but rather speech acts, the second assumption is undermined. For, on this view, the constituents of logical forms are only useful tools in guiding us to an interpretation of the utterance. There is no reason to think that an utterance's interpretation is constrained by them in the way suggested by the second assumption. However, I will not pursue this line of reply in this chapter. Nor will I here explore alternative conceptions of semantics, which seek a middle ground between the conception of semantics I have articulated, and the radical Austinean approach.

Rather, the line of response that I wish to pursue involves the denial that there is good evidence for the existence of linguistic constructions of the sort discussed by the proponents of truth-conditional pragmatics. According to the view underlying this response, the effects of context on truth-conditions are indeed limited to resolving ambiguity and providing the values to constituents of the logical forms of uttered sentences.

It is often assumed that the objects of semantic interpretation, that is syntactic logical forms, are free of lexical and structural ambiguity.<sup>11</sup> However, sometimes the sounds we hear suffer from such ambiguity. One role context plays is in helping us to decide which logical form is the one that has been uttered. That is, we draw upon extra-linguistic context to help us decide what to interpret. This is the grammatical role of context. The fact that context has a grammatical role is uncontentious and unthreatening.

The grammatical role of context solves the equation:

$$\text{utterance} + X = \text{logical form}$$

However, there are of course other roles context plays in interpretation. For example, extra-linguistic context also solves the equation:

$$\text{logical form} + \text{meaning assignments} + X = \text{truth-conditions.}$$

This is the *truth-conditional role of context*.

<sup>10</sup> For one discussion of the issues at stake, see the discussion of decoding processes versus inferential processes in chs. 1 and 2 of Sperber and Wilson (1986).

<sup>11</sup> This assumption is challenged in so-called “under-specification” approaches (cf. the essays in van Deeter and Peters (1996); cf. also the discussion of the “Scope Principle” in May (1985)).



For example, the truth-conditions of the sentence “I am a philosopher” vary from context to context. But we do not wish to count this as either a case of lexical or structural ambiguity. Rather, it is a case of indexicality. There is a broad use and a narrow use of the term “indexical”, and hence also of “indexicality”. Broadly construed, an indexical is any contextual parameter, by which I mean any primitive expression whose denotation is supplied entirely by context, perhaps guided by a linguistic rule.<sup>12</sup> In the narrow sense of “indexical”, an indexical is a proper subset of context-sensitive expressions, one with the characteristics shared by words such as “I”, “here”, and “now”, but not by “this”, “that”, “she”, and “he”. One role context plays in the determination of truth-conditions is in the assignment of values to context-dependent primitive expressions, typically unambiguous expressions with impoverished linguistic meanings. In the broad sense of “indexicality”, this is the role of context in resolving indexicality.

According to the truth-conditional pragmatist, there are truth-conditional roles of context other than the resolution of indexicality, broadly construed. If so, then not all truth-conditional effects of context are traceable to logical form. According to the second response to the truth-conditional pragmatist, there is no good reason for thinking that there are any truth-conditional roles of context aside from the resolution of indexicality, broadly construed.

My own view of the truth-conditional role of context is very conservative. First, there are expressions which are obviously indexicals in the narrow sense of the term, words such as “I”, “here”, “you”, “now”, and their brethren. Secondly, there are expressions which are obviously demonstratives, such as “this” and “that”. Third, there are expressions that are obviously pronouns, such as “he” and “she”. Overt expressions that are in none of these classes are not context-dependent. If the truth-conditions of constructions containing them are affected by extra-linguistic context, this context dependence must be traced to the presence of an obvious indexical, demonstrative, or pronominal expression at logical form, or to a structural position in logical form that is occupied by a covert variable.<sup>13</sup>

<sup>12</sup> I am using “expression” here also in a broad sense. As I use it here, it includes, for example, covert variables.

<sup>13</sup> I will assume, in this chapter, a traditional syntax involving variables. However, there is an alternative conception of syntax and semantics, in which variables are eliminated in favor of operators. Such frameworks have been recently advanced by some linguists, who claim that it has methodological advantages over frameworks involving variables (e.g., Szabolcsi (1989)). I am not quite sure how the adoption of a variable free framework would affect the discussion in this paper. Many of the principles and theses of this chapter would have to be reformulated. However, though I will not argue the point here, I do not in the end think that these reformulations would affect either the substance of my claims or the soundness of my arguments. The reason for my optimism is that the explicit variable free frameworks with which I am familiar (e.g., Cresswell (1996)) replace variables in the syntax by operators. Readings of a sentence which, on a framework with variables, may involve just postulating

If this view is correct, then any contextual effect on truth-conditions that is not traceable to an indexical, pronoun, or demonstrative in the narrow sense must be traceable to a structural position occupied by a variable. Claims of unobvious context-dependence must therefore be accompanied by arguments for the existence of a corresponding formative in the logical form of the relevant constructions.<sup>14</sup> Therefore, in this chapter, I will explore the second line of response against the proponents of truth-conditional pragmatics. That is, I will argue that we have been given no reason to abandon the thesis that the only truth-conditional role of context is the resolution of indexicality, broadly construed. If so, then we have been given no reason to rethink the view that semantics is about truth-conditions.

There are two sorts of cases I consider. The first are alleged cases of “non-sentential assertion”; utterances of expressions that do not appear to have sentential structure, yet appear to express full-blown propositions. The second concern utterances of expressions with sentential structure, which appear to express full-blown propositions, propositions that contain constituents which do not appear to be the values of any constituent in the logical form of the expression uttered.

## II

The first set of examples involves non-sentential discourse. The linguistic importance of such constructions has been emphasized by several authors (cf. Yanofsky (1978), Barton (1990), and, more recently, Staintain (1994), (1995), (1997), (1998)). By a sentence, I will mean an expression with clausal structure, containing at least a noun phrase and a verb phrase, corresponding to the traditional grammatical categories

one variable and assigning it a value (e.g., a “free” reading), sometimes involve the postulation of several operators in the syntax to be adequately reflected in a variable-free framework. Thus, where I would postulate a variable to account for such readings, Cresswell would postulate a string of operators in the syntax. My view can certainly be restated in these terms.

<sup>14</sup> Rizzi (1986) has argued that there are languages that allow understood elements that are not syntactically represented; in fact, English is one such language. The sort of examples that Rizzi has in mind are sentences such as “The sign cautions against driving over 30 m.p.h” and “John ate”. An utterance of the former expresses the proposition that the sign cautions everyone against driving over 30 m.p.h., and an utterance of the latter expresses the proposition that John ate something. However, such examples do not threaten the thesis that any contextually provided element has to be syntactically represented. For Rizzi’s examples are best understood as cases of (limited) ambiguity. What they show is that certain English verbs allow for limited type-shifting between relational, existential, and universal meanings. So, in English, “ate” is ambiguous between a two-place relation, and the result of existentially quantifying the second argument place. Similarly, “cautions” is ambiguous between a two-place relation and the result of universally quantifying the second argument place.

of subject and predicate. Call an utterance *unembedded* if and only if it is an utterance of a non-sentential expression, and it is not part of an utterance of a sentence in which that expression occurs as a constituent. So, an utterance in English of the word “water” alone (not within the context of a sentence) is an unembedded utterance. Call an utterance a *non-sentential assertion* if and only if it is an unembedded utterance that is a successful linguistic assertion. In this section, I will argue that there are no clear examples of non-sentential assertions.

If there are non-sentential assertions, then context plays more truth-conditional roles than the resolution of indexicality, broadly construed. In a non-sentential assertion, the semantic values of the words uttered, relative to that context, only make up one part of the proposition thereby expressed. Context supplies the other constituents of the proposition expressed. But then context supplies constituents to propositions expressed in assertions not merely by assigning values to constituents of the expression uttered. Rather, it provides them directly to what is asserted. But this is a truth-conditional role of context distinct from the resolution of indexicality, broadly construed.

One might have thought that the claim that there are non-sentential assertions is fairly easy to establish. After all, there are many natural languages in which there appear to be clear, fully grammatical utterances of single words, which are taken to be assertions. Such is the case in Spanish, for instance, where “corre” can be used to assert the proposition that some contextually salient man runs. It might appear that such uses do not involve the utterance of an expression with sentential structure. However, according to recent syntactic theory, there can be no subjectless sentences. What appear to be subjectless sentences in natural language actually involve covert elements in their subject positions. This, at any rate, is the import of the Extended Projection Principle of Chomsky’s Government and Binding Theory. Hence, utterances of “corre” in Spanish are not unembedded; their true structure involves a covert pronoun like element occupying the subject position. They are therefore not non-sentential assertions.

A similarly misleading set of examples involve cases like the following. Suppose that John asks:

(1) Who bought the bottle?

and Sarah responds by uttering

(2) Bill

In this case, Sarah’s utterance may appear to be a non-sentential assertion. But it is not. Rather, it is a case of syntactic ellipsis. The proposition expressed by Sarah’s utterance of (2) is plausibly taken to be the proposition that Bill bought the bottle. But

the reason it is plausible to associate this proposition with Sarah's utterance is that it is plausible to maintain that the logical form uttered by Sarah actually contained the words "bought the bottle", only covertly. This is a case of syntactic ellipsis.

It should be noted at the outset that there are some theorists who would deny that (2) in this context is a case of syntactic ellipsis. Such theorists reject the existence of real syntactic ellipsis. My arguments in what follows unfortunately do not address such theorists. I assume, as do authors who stress the importance of non-sentential discourse such as Yanofsky, Barton, and Stainton, that syntactic ellipsis is the correct theoretical account of certain ordinary linguistic phenomena. Responding to those who deny that there are any real cases of syntactic ellipsis is a foundational challenge for another time.

It is not very clear whether Yanofsky (1978) and Barton (1990) wish to establish that much of what is apparently non-sentential speech in fact consists of assertions; their explicit goal is rather to argue that not all such cases involve syntactic ellipsis, and (in the case of Barton), to supplant the syntactic ellipsis account with a novel pragmatic account of how we process non-sentential discourse. However, it is very clear that this is Stainton's desired conclusion, and it is the conclusion that is of concern in this section. In all of his papers, the way Stainton argues for the existence of non-sentential assertions is as follows. First, he produces a barrage of alleged examples of non-sentential assertions. Then, he considers a series of approaches to all of the examples he discusses. Each approach is general, in that it treats all alleged examples of non-sentential assertion in the same way. For example, one approach is to try to assimilate all alleged examples of non-sentential assertion to the case of Spanish utterances of "corre". Another approach is to try to assimilate all alleged examples of non-sentential assertion to the case of syntactic ellipsis. Finally, he rejects each strategy as at most adequate for some of the alleged examples.

The persuasiveness of Stainton's arguments is due in part to the tacit assumption that all alleged examples of non-sentential assertion must be treated by the same general strategy. However, there is no reason to accept this assumption. I do not believe that there is a uniform phenomenon underlying all apparent examples of non-sentential assertion. Many, on closer inspection, turn out to be cases of ellipsis. Others turn out not to be cases of linguistic assertion at all. Once the various examples are placed in their distinct categories, we are left without a single unproblematic example of a non-sentential assertion. Or so I will argue.

The central argument that apparent cases of non-sentential assertion differ from genuine cases of syntactic ellipsis relies principally on the fact that elliptical expressions cannot appear in discourse initial position; this is the argument strategy of both Yanofsky (1978) and Barton (1990) (cf. ch. 2), and Stainton follows them in this regard (cf. Stainton (1997, 63 ff.), (1998, 323 ff.)). According to these authors, many apparent

cases of non-sentential assertion can appear in discourse initial position. If so, they are not cases of syntactic ellipsis.

Here is an example of this sort of argument. Consider the following discourse:

- (3) (a) Bill will bungee-jump.  
 (b) John won't.

The second sentence in this discourse is a standard example of syntactic ellipsis. Now consider simply an utterance, at the beginning of a discourse, of:

- (4) John won't.

Such an utterance seems unacceptable.

Now consider the following context. Suppose Bill walks into a room in which a woman in the corner is attracting an undue amount of attention. Turning quizzically to John, he arches his eyebrow and gestures towards the woman. John replies:

- (5) a world famous topologist.

John has just uttered a phrase in isolation. It appears that John's utterance cannot be elliptical, since it occurred in a discourse-initial position. Therefore, it is implausible to assimilate cases such as utterances of (5) to syntactic ellipsis.

This sort of argument forms the backbone of the thesis that most apparent cases of non-sentential assertion are not syntactic ellipsis. However, such arguments are seriously flawed. For only in an implausibly expansive sense of "discourse-initial" does John's utterance of (5) count as discourse-initial. It is true that syntactically elliptical sentences cannot felicitously occur in the absence of a linguistic antecedent. But explicitly providing a linguistic antecedent by mentioning it is only the simplest way to provide it. There are other methods of raising linguistic expressions to salience in a conversation without explicitly using them.

For example, suppose that a group of friends, including John and Bill, has gone bungee-jumping. Every member of the group is watching Bill, who is the first to muster the courage to bungee-jump. As Bill is standing eight stories above the water on the platform of a crane, ready to plummet into the water below, Sarah, aware of John's terror of heights, turns to one of the other friends and utters (4), shaking her head. Sarah's utterance is perfectly felicitous. But it would be wrong to conclude from this that explicitly elliptical expressions can occur without linguistic antecedents. In this case, the expression "bungee-jump" has been made salient by the utterance context, and can serve as a linguistic antecedent for the syntactic ellipsis.<sup>15</sup>

<sup>15</sup> That a linguistic expression can be made salient in a context without being explicitly mentioned is neither a new nor a radical claim. For example, in the literature on E-type anaphora, it is standardly

Given that linguistic expressions can be made salient in the context in other ways than by explicitly mentioning them, we need to be careful about the use of the expression “discourse-initial”. Constructions that require explicit linguistic antecedents, such as those involving syntactic ellipsis, can occur in contexts in which the linguistic antecedents have not been used, but have been made salient in other ways. This does not show that such expressions can be used discourse-initially. It is considerably easier to make an expression salient by using it, but, with a sufficient amount of contextual cues, an expression can be made salient without using it. In evaluating the above sort of argument, we need to be certain that the alleged examples of non-sentential assertions can be used discourse-initially, where this means, felicitously used in an absolutely novel context, one in which we are assured that context has not raised any linguistic expression to salience.

To focus matters, let us consider an utterly standard example of a sentence that can be felicitously used at the beginning of a discourse:

(6) A man was walking through New York City.

(6) can be used felicitously *even if no background context has been set up at all*. This is the notion of “discourse-initial” that is common in the linguistics literature, the one that is familiar from, say, Discourse Representation Theory. A construction can occur discourse-initially just in case it can occur with minimal previous background context, whether linguistic or non-linguistic. Obviously, such a notion of discourse-initiality differs radically from the implausibly expansive one at work in the situation in the argument involving (5).

Now, it is true that syntactic ellipsis requires background context. Therefore, constructions involving syntactic ellipsis cannot occur discourse-initially, in the standard sense of discourse-initiality. However, neither can most apparent examples of non-sentential assertions. For example, a discourse-initial utterance of (5) is completely infelicitous in the standard sense of “discourse-initial”. For an utterance of (5) to be felicitous, a large amount of background context needs to be provided.

assumed that the differing acceptability of (a) and (b) demonstrate that E-type anaphora requires a linguistic antecedent:

- (a) John has a wife and she hates him.
- (b) ?John is married and she hates him.

(The minimal pair is due to Gareth Evans (cf. 1985, 147).) However, everyone is aware that, given sufficient linguistic context, (b) is perfectly acceptable. The conclusion to draw is not that E-type anaphora does not require a linguistic antecedent; the fact that (a) is always acceptable, and (b) is often not, by itself shows that it does. Rather, the natural conclusion is that an expression that requires a linguistic antecedent can be provided one by extra-linguistic context, though context needs to work hard to do so.

So, both constructions involving syntactic ellipsis and most apparent examples of non-sentential assertions cannot occur discourse-initially.

Furthermore, there is good reason to think that the background context required to license an utterance of (5) raises linguistic expressions to salience that can serve as antecedents for ellipsis. Recall the original example of syntactic ellipsis, Sarah's utterance of (2). In this case, the ellided material came from an explicit question. It is very common to respond to explicit questions by uttering what appear to be single words. However, it is plausible to suppose that the ellided material is the "standard answer schema" for the relevant question (cf. sect. 2 of Higginbotham (1993)). For example, the standard answer schema for the question (1) would be:

(7)  $\alpha$  bought the bottle.

Similarly, in the case of the world famous topologist, it is plausible to suppose that extra-linguistic context, such as Bill's gesture, and his quizzical glance at John, gave rise to the implicit question:

(8) Who is she?

John's utterance of (5) is then elliptical for "she is a world famous topologist" for the very same reason that Sarah's utterance of (2) is elliptical for:

(9) Bill bought the bottle.

I suspect that a great many apparent cases of non-sentential assertions are simply answers to implicit questions in the utterance context.<sup>16</sup> If so, then they are sentential after all.<sup>17</sup> However, many examples of apparent non-sentential assertions are clearly not cases of syntactic ellipsis. For some such cases occur discourse-initially, in the standard sense of that phrase. Consider, for example, a thirsty man who staggers up to a street vendor and utters:

(10) water

Clearly, this utterance occurs discourse-initially in every sense. However, in this case, I doubt that the thirsty man has made a linguistic speech act.

Here are two reasons to doubt that the case of the thirsty man involves a linguistic speech act. First, linguistic speech acts must determinately be made with the relevant

<sup>16</sup> For example, consider this case of Stainton's (1997, 72): "It's fair to assume that 'Potato Digging'—the bare phrase, that is—could be used on its own: You might look quizzically at a pair of mud covered boys, out in a field. I could explain their sorry state by saying 'Potato Digging. All morning.'" It is clear, in this case, that the quizzical glance gives rise to the implicit question, "What have they been doing?", to which the answer is "they have been potato digging".

<sup>17</sup> A version of the strategy I have pursued here is also defended in Fiengo and May (1996, 139 ff.). Indeed, as Fiengo and May put the moral of their discussion, "... verbalization is only tangentially related to the representations which underlie speakers' utterances".

sort of force. That is, for an act to count as a speech act of kind *k*, it must determinately be performed with the force appropriate to acts of kind *k*. For example, if the thirsty-man's utterance of (10) is an assertion, then it must be determinately made with assertoric force. However, I doubt that, in the case of the thirsty-man's utterance of (10), it is determinate that there is assertoric force. It would be equally consistent with the thirsty-man's intentions to suppose that the utterance was a request, or a command. That is, it is indeterminate what the force is with which (10) is uttered. It is therefore not a linguistic assertion, and indeed is not a genuine linguistic speech act.

Here is the second rather more complex reason why I do not believe it to be plausible that the case of the thirsty man is a linguistic speech act. Linguistic speech acts must not just be determinately made with the relevant sort of force. They also must express determinate contents.<sup>18</sup> And certainly, in the case of the thirsty-man's utterance of (10), there is no determinate content associated with the speech act. Suppose, for the sake of argument, that the speech act is an assertion. Then, the relevant sort of content is a proposition. But what proposition has thereby been expressed? The point is particularly acute if we assume that propositions are structured. Is the proposition thereby expressed the proposition that the thirsty man wants water? Is it the proposition that the vendor should give the thirsty man water? The available facts simply do not determine a determinate propositional content for the alleged assertion. And when a communicative act lacks a determinate content, it is not a linguistic speech act.<sup>19</sup>

Of course, if this last point is correct, some communicative acts involving the use of language will not count as genuine linguistic speech acts. But this is to be expected. Ordinary discourse often involves the use of complex expressions which would be counted as ungrammatical even by the utterer's own lights. For example, some people regularly start a new sentence halfway through an utterance of another

<sup>18</sup> There are several places in the literature in which this claim is challenged, e.g. in the account of incomplete definite descriptions given in Blackburn (1988). However, I do not find his account compelling; there are better accounts of the context-dependence of quantification which are consistent with this claim (see Chapter 2). The claim is not in conflict with the view advanced by Perry (1997), according to which an utterance is associated with a variety of truth-conditions. For Perry selects one notion of content to be the "official" notion of content (what he calls "content<sub>c</sub>", cf. 601), which is essentially the notion of content assumed here.

<sup>19</sup> There are analyses of vagueness according to which sentences relative to contexts containing vague terms do not express unique propositions, but rather express sets of propositions. On such accounts, "That is a heap", pointing to a heap *h*, relative to a certain context, expresses the set of propositions  $\langle h \text{ is an } F \rangle$ , where "F" is a schematic letter replaceable by non-vague heap predicates. This analysis of vagueness is simply not in tension with the claim that linguistic speech acts have determinate contents. In the case of a vague utterance, all of the different propositions in the set are structurally isomorphic; they differ only in containing different precise properties corresponding to the occurrence of the vague predicate. To account for this, it is sufficient to modify the claim in the following manner. Linguistic speech acts must express determinate contents. If a speech act expresses a set of propositions, the different propositions must be structurally isomorphic.



sentence. Such discourse involves few sentences that the utterers themselves would classify as grammatical. It is absurd to suppose that we should count such discourse as grammatical, and thereby modify syntactic theory to account for it, and this despite its (statistically speaking) relative normalcy. It is just as absurd to suppose that our conception of semantics should be modified to account for every communicative action which involves the use of language.

To say that non-sentential utterances are not linguistic speech acts, and hence not within the proper domain of study for syntax and semantics, is not to deny that they occur, or even that they are often used as vehicles of communication. A kick under the table, a tap on the shoulder, or a frown are all frequently occurring communicative actions. Indeed, one can communicate something by saying nothing at all. There is no doubt much of interest to be said about how general knowledge is brought to bear in interpreting communicative interchanges of this sort. However, it would be an error to extend the domain of linguistic theory to account for them. Such interchanges lack the distinctive features associated with linguistic communication. It is not the task of linguists to explain how communication can be effected with their use, but rather the task of the psychologist interested in rationality and ordinary inference.

There is a final set of cases that can occur discourse-initially, are clearly uttered with assertoric force, and have determinate unique propositional contents. One example given by Stainton (1995, 293) is an utterance of “nice dress”, perhaps to a woman one passes by in the street. In this case, it is fairly clear that an assertion has been made, whose content is a singular proposition about the object in question, to the effect that it is a nice dress. However, it is intuitively plausible to suppose, in this case, that the speaker simply intended her utterance to be shorthand for “that is a nice dress”. It is difficult to see how any of the resources of linguistic theory could be used to show that intuition misleads in cases of this sort.

Each and every alleged example of non-sentential assertion can be classified in one of the three ways I have described. The illusion that each strategy is unsatisfactory stems from the tacit assumption that, to be satisfactory, a strategy must work for each case of an alleged non-sentential assertion. This assumption presupposes that the “phenomenon” of non-sentential assertion constitutes a natural kind. Once this presupposition has been abandoned, it is far less clear that there are any actual everyday examples of non-sentential assertion.

### III

There is a different set of cases that has been exploited by advocates of truth-conditional pragmatics to argue that the effects of context on truth-conditional

interpretation cannot be constrained by logical form. These sorts of examples involve the consideration of sentential utterances that clearly express unique propositions. However, in such cases, it appears that there are constituents of the propositions thereby expressed that do not correspond to anything in the structure of the sentence uttered. Such entities are called in the literature *unarticulated constituents*.<sup>20</sup>

Unarticulated constituents are elements supplied by context to the truth-conditions of utterances, elements which are not the semantic values of any constituents in the actual structure of natural language sentences. That is:

x is an unarticulated constituent of an utterance u iff (1) x is an element supplied by context to the truth-conditions of u, and (2) x is not the semantic value of any constituent of the logical form of the sentence uttered.

If there are any unarticulated constituents of utterances, then context plays more truth-conditional roles than the resolution of indexicality, broadly construed. However, as I argue in what follows, the standard examples motivating the existence of unarticulated constituents are not persuasive.

My target is, in each of the examples I discuss, an *unarticulated constituent analysis* of the relevant construction. An unarticulated constituent analysis of a linguistic construction is an analysis according to which uses of that construction express propositions with unarticulated constituents. In each case, I begin by considering such an analysis. Since the supposed unarticulated constituent supplied by such an analysis is not the value of anything in the sentence uttered, there should be no readings of the relevant linguistic constructions in which the unarticulated constituent varies with the values introduced by operators in the sentence uttered. Operators in a sentence can interact only with variables in the sentence that lie within their scope. But, if the constituent is unarticulated, it is not the value of any variable in the sentence. Thus, its interpretation cannot be controlled by operators in the sentence.

<sup>20</sup> Places in which the notion of an unarticulated constituent is used in this manner include Sperber and Wilson (1986, ch. 4), Recanati (1993, sect. 14.3), and Bach (1994, sect. 2). Crimmins (1992, ch. 1) also uses the notion of an unarticulated constituent against a version of the thesis that all context-dependence is traceable to structure. However Crimmins is substantially more cautious than the other advocates of unarticulated constituents. His target is not the view that all context-dependence is traceable to logical form, as I have presented this thesis, but the much more implausible view that contextual effects on truth-conditions are restricted to providing the values of expressions in the *apparent structure* of the sentence. Therefore, he should not be assimilated to my targets. A similar point does not hold of the article in which the vocabulary was introduced, Perry (1986), since, in his (1998), Perry is clear that the phenomenon of interest to him is what he calls a “truly unarticulated constituent”, which is not the value of an unpronounced item in the actual structure of a sentence (see his fn. 4).

The arguments I provide will also be sufficient, in each case, to refute what one may call a *narrow indexical* analysis of the relevant construction. In the narrow sense of the term “indexical”, it applies to words such as “I”, “here”, “you”, and “now”. The three central features of such words are: first, that they are primitive lexical items; second, that they are not bindable by operators; and, third, that their interpretation shifts from context to context. An unarticulated constituent analysis of an expression is closely related to the claim that the relevant expression is an indexical in the narrow sense of the term, a primitive lexical item whose content varies from context to context, and which is resistant to binding by a variable-binding operator with scope over it. My arguments against unarticulated constituent analyses of the constructions I discuss will also show that the relevant expressions are not indexicals.

The reason I extend my arguments to narrow indexical analyses of the constructions I discuss is not because such analyses have been proposed or defended for such constructions. It is rather because narrow indexical analyses of certain philosophical expressions, such as “true” and “knows” are common in the philosophical literature. If it can be shown that, in the non-philosophical areas of our speech, narrow indexicality is restricted to obviously indexical expressions such as “I”, “here”, and “now”, then the thesis that philosophically controversial expressions such as “true” or “knows” are narrow indexicals will thereby be placed into doubt. This is an additional benefit of the arguments of this section.

The first step in my arguments is to show that there are readings on which the interpretation of the alleged unarticulated constituent is controlled by an operator in the sentence. It follows that an unarticulated constituent analysis is incorrect. The second step is to draw the conclusion that there is in fact a variable in the logical form of the sentence uttered, whose value is the contextually supplied constituent.

The first step in the argument against an unarticulated constituent analysis is sufficient to refute a narrow indexical account of the relevant construction for the following reason. Showing that the interpretation of the alleged unarticulated constituent can be controlled by a higher operator is tantamount to showing that the contextually supplied element is the value of a bindable constituent in the logical form of the relevant construction. Since indexical expressions, narrowly construed, are not bindable, it follows that the context-dependence in question is not due to the presence of indexicality, narrowly construed.

The second of the above steps requires one methodological presupposition. Though it is a little unwieldy to state in detail, it is quite innocent. Roughly, the presupposition is that, for explicit quantifier expressions, within a clause, semantic binding and syntactic binding coincide. That is, bound readings within a clause are due to the existence of a variable binding operator standing in a certain structural relationship to a co-indexed variable in that clause.

Let me make this explicit. Suppose  $\alpha$  is an explicit quantifier expression.<sup>21</sup> Let us say that  $\alpha$  *semantically binds*  $\beta$  if and only if the interpretation of  $\beta$  systematically depends upon the values introduced by  $\alpha$ . Then:

### The Binding Assumption (BA)

If  $\alpha$  and  $\beta$  are within the same clause, and  $\alpha$  semantically binds  $\beta$ , then  $\alpha$  either is, or introduces, a variable-binding operator which is co-indexed with, and stands in a certain specified structural relation to, a variable which is either identical to, or is a constituent of,  $\beta$ .<sup>22</sup>

BA is familiar from the syntax and semantics of first-order predicate logic, in which bound readings are due to the existence of variable-binding operators having co-indexed variables within their scope.<sup>23</sup> According to it, binding within a clause is fundamentally a syntactic phenomenon.<sup>24</sup>

BA is natural in a semantics involving structured propositions. Structured propositions contain objects and properties of various sorts. One natural treatment of

<sup>21</sup> Among “explicit quantifier expressions”, I include what David Lewis has called “adverbs of quantification”, such as “always”, “usually”, and “sometimes”.

<sup>22</sup> I am appealing here to a broader notion of variable than the one corresponding to the use of the term “variable” in the theory of Government and Binding. For example, a widely adopted definition of the term occurs in Koopman and Sportiche (1982/3), in which a variable is defined as any expression in an A-position that is locally A-bar bound. In the sentence “[Every woman]<sub>i</sub> t<sub>i</sub> loves her<sub>i</sub> mother”, “her<sub>i</sub>” is a variable in the broad sense, but not in this latter sense. One can define variable-hood in the broad sense in terms of the concepts of Government and Binding Theory as follows. Let the antecedence relation be that relation that holds between  $\alpha$  and  $\beta$  if and only if  $\alpha$  is the immediate antecedent of  $\beta$ .  $\gamma$  is a variable in the broad sense if and only if it stands in the weak ancestral of the antecedence relation to a variable in the technical sense of GB theory.

<sup>23</sup> This assumption is also consistent with the logical tradition that descends from the final section of Kazimierz Ajdukiewicz’s masterful (1967). According to this tradition, there is only one sort of variable-binding operator, what Ajdukiewicz, following Russell, calls the circumflex (Ajdukiewicz (1967, 227), David Lewis (1983, 211) calls a “binder”, and Max Cresswell calls a  $\lambda$ -abstractor. If one assumes obligatory syntactic quantifier raising, the assumption is sound in such a framework, because each raised quantifier expression introduces an occurrence of one of the variable-binding operators. This is also the treatment of Heim and Kratzer (1998).

<sup>24</sup> Certain special sentential expressions, such as the modal expressions “necessary” and “possible”, are often treated as expressions that semantically bind expressions without the mediation of variables. However, the class of such expressions is (or should be) highly restricted. Such a treatment of modal expressions is justified on the assumption that it is appropriate to take possible worlds as entities relative to which the truth of propositions is evaluated, rather than parts of the contents of propositions. In contrast, I do not think it is plausible to take temporal expressions as operators of this kind. First of all, times are generally assumed to be a regular part of the content of propositions. Furthermore, there are a host of independent objections to the treatment of temporal expressions as operators rather than as predicates of times. Be that as it may, the existence of a highly restricted class of expressions of this sort is consistent with the arguments that follow, as long as this class does not contain the standard quantifiers.

binding within such a framework is to suppose that quantifiers are associated with variable-binding operators. When the quantifiers undergo quantifier movement, they introduce variable-binders, which bind any variables within their scope. The effect of such variable binders is to transform open sentences into names of properties, or alternatively, propositional functions. The property named, together with the quantifier denotation, are then elements of the structured proposition expressed (cf. Salmon (1986, 157)). Such a treatment of binding in a structured proposition framework is in accord with the above assumption.

According to most semantic frameworks, including the one just discussed, bound variables do not have independent denotations. However, one might adopt a more liberal conception of structured propositions, according to which they do not contain just objects, properties, and quantifier denotations, but also contain elements which correspond to the occurrence of bound variables. Bound readings would then result from the semantic interactions between these elements and the denotations of the corresponding variable binders. Motivations for such a treatment of binding can come from a variety of sources; for example, a desire to maintain a particularly severe form of compositionality (cf. Lewis (1983, 212)), or a desire to preserve a “strong sort of semanticism about logic” (Varzi (1993)).

This framework suggests a way of avoiding commitment to BA.<sup>25</sup> Assuming such a framework, one could simply reject the thesis that bound variables must always exist in the syntax for a bound reading to occur. On this account, the semantic elements corresponding to bound variables can be supplied by the semantics, with no corresponding syntactic element denoting them.

However, given her commitments, the advocate of truth-conditional pragmatics should not accept this latter possibility. For advocates of truth-conditional pragmatics hold that each element of the proposition expressed must either be the value of some element in the syntactic structure, or provided by pragmatic mechanisms. It is easy to see how an object or a property could be provided by pragmatic mechanisms; it need only be made salient in the context either by the speaker’s intentions or contextual cues, depending upon one’s account of salience. However, denotations of bound variables are odd, theoretically complex entities. It is difficult, if not impossible, to see how, on any account of salience, such an entity could be salient in a context. Certainly, neither it, nor instances of it, could be perceptually present in the context. It is equally difficult to see how speaker intentions could determine reference to such an entity.<sup>26</sup>

<sup>25</sup> I am grateful to Jeff King for emphasizing this possibility to me.

<sup>26</sup> Jeff King (1995) advances an account of propositions according to which they contain variables—the actual linguistic entities. A variable in the sentence is taken to contribute itself to the proposition. However, King’s rather idiosyncratic conception of propositions is motivated on the thesis

An entity such as a denotation of a bound variable is a theoretical posit, part of the machinery of a particularly complex semantic theory. It is not something about which we have beliefs or intentions. They are therefore not supplied by pragmatic mechanisms. Given that such entities are not supplied by pragmatic mechanisms, then, they must be part of semantic interpretation. But, given the commitments of the advocates of truth-conditional pragmatics, if an account of binding involving such entities is adopted, they must then be the values of elements in the syntactic structure of sentences. The entities which denote them, of course, are variables. Therefore, this way of avoiding commitment to the Binding Assumption is not available to the advocate of truth-conditional pragmatics. Indeed, I suspect that the advocate of truth-conditional pragmatics is in the end committed to BA. This is not in itself worrying, of course, since standard treatments of binding are fully consistent with BA. However, as I now show, BA, together with some empirical facts, poses serious difficulties for those who believe in the existence of unarticulated constituents.

Here is an argument for the existence of unarticulated constituents, due originally to John Perry.<sup>27</sup> Consider the sentence:

(11) It's raining.

According to this argument, it is plausible that (11) contains a covert temporal variable, so that its true representation is more like:

(12) It is raining (t).

But what an utterance of (11) asserts is not just that it is raining at a certain contextually provided time. Rather, it asserts that it is raining at a certain contextually provided time at a certain contextually provided place. But surely it is implausible to posit a place variable in addition to a temporal variable. It is surely more plausible to supply the place to the truth-conditions of an utterance of (11) directly, without mediation of a variable.

that a proposition is a kind of "shadow" of the sentence that expresses it. To suppose that a sentence not containing a variable could regularly express one of King's propositions containing a variable is not to think of the sentence/proposition relation as King thinks of it. The idea that a proposition could contain a linguistic element such as a variable is only plausible given King's background account of the intimate relation between a proposition and the sentence that expresses it.

<sup>27</sup> cf. Perry (1986, sect. 1). Perry's motivation for introducing unarticulated constituents is to argue that each of us is an unarticulated constituent of our own thoughts at the level of "the most basic kind of self-knowledge" (ibid., 138). Perry's aim is thereby to justify the Humean claim that we have no representation as of ourselves. However, I do not believe that there are convincing reasons in favor of the Humean claim. The claim derives its initial appeal from an overly restrictive sense of "representation". We certainly do not have representations of ourselves that are closely analogous to our representations of entities external to our bodies that are perceived through visual or auditory means. But this does not provide evidence that we have no representation as of ourselves. It simply leads to the thesis that our representations of ourselves have special features. For a development of this line of thought, see Cassam (1997).

Informally, here are a few of the relevant details of an unarticulated constituent analysis of (11). Suppose “t” is a variable ranging over times, and “l” a variable ranging over locations. The interpretation of “rains” would then be:

Den(“rains”) relative to a context  $c =$  that function  $f$  that takes  $\langle t, l \rangle$  to True if it is raining at  $t$  and  $l$ , where  $l$  is the contextually salient location in  $c$ , takes  $\langle t, l \rangle$  to False if it is not raining at  $t$  and  $l$ , where  $l$  is the contextually salient location, and is undefined otherwise.<sup>28</sup>

According to the unarticulated constituent analysis, the structure of (11) is as in (12). Therefore, its truth-conditions would be given by a clause such as:

R: “It is raining( $t$ )” is true in a context  $c$  if and only if the denotation of “rains” takes  $\langle t, l \rangle$  to the True, where  $l$  is the contextually salient location in  $c$ .

Clause R is a standard unarticulated constituent clause. It captures the intuition that the place variable is supplied directly by context, rather than first to a variable in the logical form of (11).

However, it is incorrect. Consider the sentence:

(13) Every time John lights a cigarette, it rains.

One natural interpretation of (13) is:

(14) For every time  $t$  at which John lights a cigarette, it rains at  $t$  at the location in which John lights a cigarette at  $t$ .

The problem this example raises for an unarticulated constituent analysis is as follows. There is no way to derive this interpretation of (13) with the use of this sort of analysis. Rather, the only reading predicted by an unarticulated constituent analysis is:

(15) For every time  $t$  at which John lights a cigarette, the denotation of “rains” takes  $\langle t, l \rangle$  to the True, where  $l$  is the contextually salient location in the context of utterance of (13).<sup>29</sup>

If one postulated, in addition to the temporal variable, a variable in the logical form of the embedded sentence “it rains”, whose value is the location at which it rains, one can capture both readings. Surely, what the evidence suggests is that this account is preferable to any unarticulated constituent analysis. If so, then the location is the value of a variable in the logical form after all.

<sup>28</sup> So, the function expressed by “rains” relative to a context  $c$  is undefined for all  $\langle t, l \rangle$  such that  $l$  is not the contextually salient location in  $c$ .

<sup>29</sup> Due to pragmatic factors, (15) is not a particularly salient reading of (13). However, suppose that John is a mad scientist, who has established a connection between his cigarette lighter and a certain location  $l$ , such that whenever he lights a cigarette with it, it rains at location  $l$ . In this situation, standing at location  $l$ , (13) may be uttered with the interpretation as in (15).

There are several ways to capture these readings of (13). One is to replace the assumption that “rain” introduces a hidden temporal variable with the assumption that it introduces a hidden situation or event variable, which can either be bound, as in (14), or free, as in (15). The situation variable brings with it information about the time and place at which it occurs. Alternatively, one may suppose that when “rain” occurs in a sentence, it co-occurs with a temporal node and a locational node. Occupying the phrases are two variables, which may either occur bound, as in (14), or free, as in (15).

According to this latter account, in the logical form of (11), “rain” occurs with two open positions. Each open position is filled by a term of the form “ $f(x)$ ”.<sup>30</sup> The first function maps entities to times, and the second function maps entities to locations. Context supplies the value of the function variables “ $f$ ” and “ $g$ ”. In the usual case, context supplies the identity function to these function variables. So, in the usual case, the value of “ $x$ ” and “ $f(x)$ ” is the same, and the value of “ $y$ ” and “ $g(y)$ ” is the same. But in examples such as (13), context supplies a function different from the identity function to one of the higher-order variables. In the case of (13), the temporal node contains a complex variable “ $f(t)$ ” and the locational node contains a complex variable “ $g(t)$ ”. When (13) is evaluated with respect to a context, “ $f$ ” is assigned the identity function, and “ $g$ ” is assigned a function from times to locations. In the case of a sentence such as “Every place John goes, it rains”, “ $g$ ” is assigned the identity function, and “ $f$ ” is assigned a function from locations to times.

However, the further details of these accounts need not detain us. For our purposes, it is not important to decide between competing accounts which both involve variables; it is only important to note the failure of the unarticulated constituent account. The problem with the unarticulated constituent analysis is that it only predicts one of the two available readings for (13). In contrast, an account involving the postulation of a location variable predicts both readings. For variables can either be bound or free. An account involving variables, therefore, predicts there to be two readings of (13), one in which the value of the relevant variable is supplied by context, as in (15), and one in which it is bound, as in (14). As we have seen, this prediction is borne out by the facts.

The above considerations also generalize to undermine the narrow indexical analysis of “rain”. The standard analysis of the semantics of indexical expressions is due to David Kaplan (1989). According to it, there are two levels of semantic content. In the first instance, word types are associated with what Kaplan calls “characters”, which are functions from contexts to the second sort of semantic content, which

<sup>30</sup> Similar appeals to covert function variables which have first-order variables as arguments occur in the analysis of functional readings of wh-questions (cf. Engdahl (1986); Chierchia (1993)). Chierchia (1995, 225–7) provides evidence for the syntactic reality of such functional variables by appeal to weak crossover considerations.



Kaplan calls, simply, “content”. The content of non-indexical, demonstrative, and pronominal expressions are not sensitive to context, and so their characters are constant functions from contexts to contents. Indexicals and unbound pronouns and demonstratives do, by contrast, have a content that varies with context. As a result, the character of such an expression is a non-constant function from contexts to contents. For example, the character of “I”, on this account, is a function from contexts to contents. Given a context, it yields a constant function from possible worlds and times to the speaker in that context.

One can give a narrow indexical analysis of “rains” to account for the truth-conditional variation exhibited by different uses of (11). On this account, “rain” is an indexical expression, narrowly construed. Therefore, it is associated with a non-constant character. Given a context, the character of “rain” yields a function from possible worlds and times to truth-values. This function is the content of “rain” in that context. The content yielded by the character of “rain” relative to a context *c* is that function from worlds and times to truth-values that yields Truth if it is raining at that world at that time in the salient location of *c*, and yields False otherwise.

However, examples such as (13) show that the narrow indexical analysis of “rains” is incorrect. In cases such as (13), the location parameter is bound. But if the location parameter is a contextual parameter, as it is if the narrow indexical analysis is correct, then it is simply not accessible for binding, any more than the speaker coordinate is. Therefore, the narrow indexical analysis of “rains” is incorrect.<sup>31</sup>

Philosophers and linguists have both used comparative adjectives to motivate the notion of an unarticulated constituent (e.g. Bach (1994, 128); Heim and Kratzer (1998, 71)). Consider the sentence:

(16) Sherman is small.

The truth-conditions of (16) vary with context. Suppose Sherman is 6 feet tall. If what is at issue in the context of an utterance of (16) are professional basketball-players, then that utterance expresses a true proposition. However, in a context in which what is at issue are junior high-school students, an utterance of (16) expresses a falsehood. However, one might think, for whatever reason, that the actual structure of (16) does not contain a “hidden” variable whose value is the contextually relevant comparison class. If so, then one should seek a semantic rule which provides the comparison class “directly”, without mediation of a variable.

<sup>31</sup> As Maria Bittner pointed out to me, there is one theoretical option remaining if one wishes to deny that sentences such as (11) and (13) involve covert variables hidden in their logical form that are accessible to binding. According to this option, “rains” is itself a variable, that in a sentence such as (13) is bound. More precisely, “rains” is a pronominal expression. Pronouns are ambiguous between deictic and bound readings. In a sentence such as (11), we see the deictic reading of “rains”. In a sentence such as (13), we see the bound reading of “rains”. I will not pursue this option here.

The rule that would be required to supply the full truth-conditional interpretation of utterances of (16), on an unarticulated constituent view, would be roughly as follows:

Rule C: Den(“small”) relative to a context *c* is the set of things of size less than *s*, where *s* is the standard made salient in *c*.

However, Rule C is incorrect, and for a similar reason as clause R. The sentence:

(17) Most species have members that are small.

has the readings given in (18) and (19):

(18) Most species *S* have members that are small for *S*.

(19) Most species *S* have members whose size is below *s*, where *s* is the standard made salient by the utterance context.

Rule C only allows for the derivation of (19), and not for the equally natural (18). Therefore, an unarticulated constituent approach to comparative adjectives is incorrect. Comparison classes are the values of contextual variables correlated with comparative adjectives.

These considerations also can be generalized to refute a narrow indexical analysis of comparative adjectives. However, at this point, laying out the details is merely a formal exercise, the details of which I leave to the reader.

Another example one could give to argue for the existence of unarticulated constituents involves sentences containing quantifier expressions. Consider the sentence:

(20) Every bottle is green.

Relative to different contexts, (20) has different truth-conditions. Relative to one context, (20) could express the proposition that every bottle recently purchased by Bill is green; relative to another, the proposition that every bottle in the house is green. Thus, context supplies a property that restricts the quantification. However, one might think that there is no variable in the logical form of (20) whose value is the required property.

Here are some of the informal details of an unarticulated constituent account of quantifier domain restriction. On this account, nouns such as “bottle” are treated as denoting different sets in different contexts.<sup>32</sup> For example, the denotation of “bottle” would be given by a rule such as:

<sup>32</sup> Quantifier domains actually are better treated as more intensional entities, such as properties (see p. 102). But treating them as sets does not affect the point I am making here.

- (21) Den(“bottle”) relative to a context  $c$  = the set of bottles that are in the domain salient in the context  $c$ .

Given a rule such as (21), one can account for the differences in truth-conditions between different utterances of (20). Relative to a context in which the salient domain is the set of things in the house, (20) will express the proposition that every member of the set of bottles in the house is green, whereas relative to another context, it will express a different proposition.

However, this account of quantifier domain restriction is unsatisfactory. Consider the following sentences:

- (22) a. In most of John’s classes, he fails exactly three students.<sup>33</sup>  
 b. In every room in John’s house, he keeps every bottle in the corner.  
 c. Whatever office you go to, the supervisor is always unavailable.  
 d. Whatever John does, most people turn up late for the experiment.<sup>34</sup>

In none of these cases does the unarticulated constituent analysis of quantifier domain restriction yield the correct result. Consider (22a). The natural interpretation of this sentence is:

- (23) In most of John’s classes  $x$ , he fails exactly three students in  $x$ .

However, the unarticulated constituent analysis only predicts the absurd reading:

- (24) In most of John’s classes  $x$ , he fails exactly three students in the domain salient in the context of utterance of (22a).

The unarticulated constituent analysis of quantifier domain restriction is therefore unsatisfactory.<sup>35</sup>

<sup>33</sup> This sort of example is discussed at length in von Stechow (1994, sect. 2.2.2).

<sup>34</sup> These last two examples are from Cooper (1996). Cooper uses these examples to argue for the existence of bound resource situation variables. However, I do not think that the situation semantic treatment of quantifier domain restriction is satisfactory, essentially for the reasons given in Soames (1986).

<sup>35</sup> Kent Bach suggested to me the possibility that in (22a) and (22b) the initial prepositional phrase has undergone movement from a structure such as:

- (a’) He fails exactly three students in most of John’s classes.

However, this analysis is easily refuted. (a’) is ungrammatical if “he” and “John” are co-indexed. The explanation for its ungrammaticality is that it is a violation of Principle C of the Binding Theory. If (22a) were derived from (a’) via movement, we would therefore expect a strong crossover violation in (22a). But (22a) is perfectly grammatical. Therefore, (22a) is not derived from (a’). A similar point holds for (22b). Furthermore, no similar account is even remotely possible in the case of (22c) and (22d).

A similar point holds for the examples (22b–d). In each case, the domain of the second quantified expression varies with the values introduced by the initial quantifier expression. Therefore, what we have been assuming about the relation between semantic binding and syntactic binding outlined above, what this evidence demonstrates is that there are bindable variables in the logical form of sentences containing quantifier expressions whose values are quantifier domains.

Another favorite example of unarticulated constituents comes from “relational expressions”, such as “home”, “enemy”, or “local” (e.g. Crimmins (1992, 151); Bach (1994)). Consider the sentence:

(25) David is at home.

What an utterance of (25) expresses is the proposition that David is at the home of N, where N is a contextually salient person (possibly David himself). Similarly:

(26) John visited a local bar.

(27) Bob faced an enemy.

express, respectively, the proposition that John visited a bar that is local to N, where N is a contextually salient person, and the proposition that Bob faced an enemy of N, where N is a contextually salient person. However, one might think that it is implausible to postulate variables in the logical form of these sentences whose values, relative to contexts, are contextually salient persons.

Here is an interpretation for “home” that would provide the contextually salient person as an unarticulated constituent:

(28) Den (“home”) relative to  $c$  = the home(s) of N, where N is the contextually salient person in  $c$ .

Let us suppose that there is a temporal variable in the logical form of a sentence such as (25). Then, the truth-conditions of (25), on an unarticulated constituent view, are as follows:

“ $x$  is at home ( $t$ )” is true in a context  $c$  if and  $x$  is at the home of N at  $t$ , where N is the contextually salient person in  $c$ .

However, (28) is an incorrect interpretation rule. Consider the sentence:

(29) Everyone is at home.

Ignoring the context-sensitivity of “everyone”, (29) has two possible interpretations:

(30) Everyone  $x$  is at the home of  $x$ .

(31) Everyone  $x$  is at the home of N, where N is the person made salient by the utterance context.

However, (28) only allows for the derivation of (31), and not for the equally natural (30). Therefore, (28) is an incorrect interpretation rule. The word “home” is accompanied in logical form by a contextual variable which is accessible to binding by a higher operator.

A similar point holds for other relational expressions. For example:

(32) Every newspaper reporter went to a local bar to hear the news.

has, among its different readings, the one given in (33):

(33) Every newspaper reporter  $x$  is such that  $x$  went to a bar local to  $x$  to hear the news.

Similarly:

(34) Every warrior faced an enemy.

has the reading given in (35):

(35) Every warrior  $x$  faced an enemy of  $x$ .

None of these readings would be available if the contextually supplied elements relevant for the truth-conditions of (26) and (27) were unarticulated constituents.<sup>36</sup>

It is also worth emphasizing that there is other good syntactic evidence for the existence of variables in constructions involving relational expressions. For example, such constructions give rise to weak crossover effects.<sup>37</sup> What this shows is that the variable element in relational expressions has the syntactic properties of explicit pronouns. Consider, for example, the following minimal pairs. In each of them, the relational expression has the same binding properties as the corresponding explicit pronoun:

- (36) a. \*Her<sub>i</sub> local bar sponsored [every reporter]<sub>i</sub>.  
 b. \*A local bar sponsored every reporter. (where the bar is the reporter’s local bar)
- (37) a. ? Her<sub>i</sub> trip home made [every reporter]<sub>i</sub> nervous.  
 b. ? The trip home made every reporter nervous. (where the home is the reporter’s home).

<sup>36</sup> For extensive discussion of relational expressions, see Mitchell (1986) and Partee (1989).

<sup>37</sup> Thanks to discussion here with Jim Higginbotham, who also emphasizes this point in his (ms). Higginbotham’s purpose differs from mine, in that he does not use the binding facts to argue for the existence of explicit formatives. Rather, he is operating with a more abstract conception of syntactic representation than the one at work in this chapter (cf. Williams (1995) for a similar conception of syntax).

Similarly, in generic contexts, weak crossover is relaxed both for explicit pronouns and for the variable element in relational expressions:

- (38) a. [Her]<sub>i</sub> trip home makes [every reporter]<sub>i</sub> nervous.  
 b. The trip home makes every reporter nervous. (where the home is the reporter's home)

This evidence strongly suggests the existence of a covert pronominal element in relational expressions.<sup>38</sup>

The sorts of arguments I have given above generalize to a host of other cases. For example, consider:

- (39) There is enough beer in the house.

The truth-conditions of an utterance of (39) depend on context.<sup>39</sup> If there are twenty bottles of beer in the house, then there is enough beer in the house for a small dinner party, but not a raucous gathering. However, higher operators can control the interpretation of the context dependent element, as in:

- (40) a. There is always (usually/sometimes) enough beer in the house.  
 b. Whenever John visits, there is not enough beer in the house.

Several other cases of this sort are discussed in Cresswell's important (1996), albeit in his variable free framework.<sup>40</sup> In all of these cases, an unarticulated constituent analysis is not tenable.

Here is a possible response to the above arguments. I present the response with the argument involving "home", though it should be clear how it generalizes to the other arguments I have given. Consider again:

- (25) David is at home.  
 (29) Everyone is at home.

<sup>38</sup> Peter Culicover and Ray Jackendoff challenge this sort of argument in their (1995). In particular, they first argue (sect. 2.6) that the element with a variable interpretation in "something else" does not obey Principle C of the Binding Theory. Then, they assume without argument that all other implicit arguments pattern like "something else" (sect. 4.1). However, their evidence that the variable element in "something else" does not obey Principle C is weak, and is equally evidence for the hypothesis that "something else" has an underlying syntactic structure similar to that of "something other than  $\alpha$ ". They consider this objection, but misconstrue it as the implausible proposal that "something else" is to be reconstructed as "something other than  $\alpha$ " at logical form. This is simply not the objection. On the natural view, what the evidence shows is that "something else" already has the syntactic structure of "something other than  $\alpha$ ", due to the presence of covert empty elements. No reconstruction is needed.

<sup>39</sup> I am grateful to Delia Graff for supplying the example.

<sup>40</sup> For examples involving modal accessibility relations, see Cresswell (1996, 56–7). For examples involving degrees of comparison, see *ibid.* (59–60).

According to this response, the occurrence of “home” in (25) is a different *word* from the occurrence of “home” in (29). What is phonetically realized as “home” is in fact ambiguous. In (25), “home” does not have an argument place for contextually salient individuals. The truth-conditionally relevant entity is added via an unarticulated constituent rule in the semantics or the pragmatics. In (29), by contrast, “home” does have an argument place for individuals, that is bound by the quantifier “everyone”. If so, then the argument I have given does not show that the occurrence “home” in sentences such as (25) brings with it a variable whose value is supplied by context. Rather, all it shows is that there is a phonetically similar word which bring such a variable with it.

However, this response is unsatisfactory. Consider the following discourse:

(41) David is at home. In fact, everyone is.

There are two interpretations of the second sentence in (41):

- (42) a. everyone  $x$  is at  $x$ 's home.  
 b. everyone  $x$  is at the home of  $N$ , where  $N$  is the contextually salient person in the utterance context of (41).

If the response we are considering were correct, (42b) would not be an available reading at all.

Here is why reading (42b) would not then be available. The second sentence in (41) is a case of syntactic ellipsis. According to standard theories of ellipsis, the material following the copula “is” in the first sentence of (41) is either copied or reconstructed in the logical form of the second sentence.<sup>41</sup> If the response we are considering is correct, then the predicate in the first sentence of (41) would not contain a variable, and so the occurrence of “at home” in the logical form of the second sentence of (41) would then also not contain a bindable variable. Thus, the second sentence of (41) would not permit a bound reading of a variable, as in (42b), since there would be no variable there to bind. But it does. Therefore, the first sentence of (41) contains a variable of the relevant sort after all.

Of course, the distinction between the readings in (42a) and (42b) is just the familiar “strict/sloppy” dichotomy found whenever overt pronouns interact with ellipsis. For example, the sentence:

(43) John likes his brother, and Bill does too.

is ambiguous between:

<sup>41</sup> The argument to follow does not depend upon a copy theory of ellipsis; it would work equally well under the assumption that ellipsis amounts to Phonological Form deletion under a parallelism requirement, as suggested in Chomsky (1995, 125 ff.).

- (44) a. John likes John's brother, and Bill likes John's brother.  
 b. John likes John's brother, and Bill likes Bill's brother.

The standard theoretical account of the distinction between these two readings is that, in the first case, the "strict" reading, the pronoun "his" is free, whereas in the second case, the "sloppy" reading, the pronoun "his" is bound. The fact that there are strict/sloppy ambiguities in ellipsis involving relational expressions is fully explained by the supposition that there are pronominal elements in these constructions. Furthermore, if there are no such elements, the existence of strict/sloppy ambiguities is left unexplained. Therefore, the existence of such ambiguities is powerful additional evidence for the existence of a pronominal element in relational expressions.

Furthermore, this dialectic generalizes to every construction we have discussed so far. Sentences such as

- (45) Bill dislikes three people. John does too.

demonstrate that there are strict/sloppy ambiguities in the case of quantifier domain restriction. It is clear that there is a "strict" reading of (45), where the quantifier domain restriction for "three people" is the same in the first clause as it is in the ellided one. However, there is also a sloppy reading of such constructions. Suppose that someone is arguing that many men have troubled relations with their families, and is using Bill and John as evidence. With respect to such a context, (45) can express the proposition that Bill dislikes three people in Bill's family, and John dislikes three people in John's family.

Similarly, consider:

- (46) John is too old. Jill is too.

It is clear that there is a strict reading of (46). However, there is also a sloppy reading. Suppose that John, a 42-year-old professional swimmer and Jill, a 23-year-old professional gymnast, have decided to wed. Shocked at their age difference, I ask Bill how John and Jill can relate to one another, to which he replies by uttering (46). Relative to this context, (46) can express the proposition that John is too old for his sport, and Jill is too old for hers. Thus, in all of the examples we have discussed, we see behavior that is best explained by the postulation of a covert pronominal element.

A second objection to my arguments is as follows.<sup>42</sup> Consider again:

- (25) David is at home.

According to the unarticulated constituent account, what the semantics assigns to (25) is a property, which is then "enriched" into a proposition. My arguments have

<sup>42</sup> I owe this objection to an anonymous referee.



shown that there is a variable associated with “home” in constructions such as (25), which is required to account for constructions such as (29). However, the defender of unarticulated constituents may maintain that this is not all the hidden syntactic structure associated with (25). Rather, she may say that the true syntactic structure of (25) is:

(47)  $\lambda x(\text{David is at home } x)$

In this way, the defender of unarticulated constituents can both maintain her thesis that the semantic content of (25) relative to a context is a property, rather than a proposition, and account for sentences such as (29).

However, this objection is simply not open to the advocate of truth-conditional pragmatics. The argument that the semantic content of (25) is a property rests on the thesis that it is always illegitimate to postulate structure on semantic grounds. It is thereby deemed illegitimate to postulate a variable in the syntactic structure of (25) on purely semantic grounds. But this objection requires the postulation, not just of a variable, but *also* of a hidden lambda-abtractor, all in the service of rescuing the semantic thesis advocated by the defender of truth-conditional pragmatics. The objection is therefore inconsistent with the justification for truth-conditional pragmatics.

Furthermore, there are positive reasons to reject this proposal. Consider:

(48) David is at home. Bill is too.

Consider the reading of the first sentence of (48) in which David is at his own home. In this case, there are two readings of the second sentence of (48):

- (49) a. Bill is at Bill’s home.  
 b. Bill is at David’s home.

According to the standard explanation of this ambiguity, the distinction between these two readings is due to whether a variable in the ellided constituent is bound or free. In this case, the standard explanation would account for the ambiguity by the hypothesis that when the variable element in “at home” is controlled by “David”, we obtain reading (49a), and when it is free, and assigned David by the context, we obtain reading (49b). However, if the true logical form of the first sentence of (48) were (47), then the variable in the ellided constituent would always be bound by the lambda-abtractor. We should therefore not expect reading (49b). The presence of strict readings in constructions of this sort therefore provides a decisive refutation of this proposal.

Here is a final objection to my arguments. In two discussions of relational expressions, Partee (1989) and Culicover and Jackendoff (1995), the existence of bindable interpretations for relational expressions is discussed at length, but an

account of the phenomena that posits empty elements in syntactic logical forms is rejected. However, the rejection of the account in terms of empty elements is based upon the rejection of a premise that has a high degree of plausibility. The premise in question is that the objects of semantic interpretation are syntactic logical forms, where these are understood as the final representations produced by the best syntactic theory. Partee rejects an account of the phenomena in terms of empty elements in actual logical form, because she thinks that what semantic theory interprets are discourse representation structures (DRSs), and it is on this representational level that the phenomenon of bound relational expressions is explained. Culicover and Jackendoff reject an account in terms of empty elements in actual logical form, because they hold that the object of semantic interpretation are what they call “conceptual structures” (CSs). The phenomenon of bound relational expressions is to be explained, according to them, by the existence of empty elements in conceptual structure. According to these theorists, the binding of relational expressions is not to be explained by the existence of empty elements in standard syntactic structures, but rather via the existence of empty elements in alternative formal levels of representation.

There are two ways of understanding such claims. According to the first, the suggestion is that syntactic theory produces syntactic logical forms, which are then jettisoned in favor of other structures, which are the input to semantic interpretation. In the case of Partee, these second structures are DRSs, and in the case of Culicover and Jackendoff, they are CSs. This picture of interpretation is *prima facie* difficult to accept. According to it, the interpretative process involves the production of an interpretively superfluous level of representation, namely the output of the syntactic mechanism. We would need a massive amount of empirical and methodological motivation to justify the added complexity such an interpretive process involves over straightforwardly applying a semantic interpretation to the output of our best syntactic theory. An evaluation of our grounds for this added complexity goes well beyond the scope of this chapter.

Though claims of the sort made by Partee and Culicover and Jackendoff are usually presented in the first manner, they are perhaps better understood in a second way. Both Discourse Representation Structures and Conceptual Structures are syntactic levels of representation, themselves in need of interpretation. Another way to construe Partee’s suggestion is that Discourse Representation Structure is the correct syntactic representation of natural language sentences; *mutatis mutandis* for Culicover and Jackendoff and CS. On this reading, what all these authors reject is the thesis that the bindability of relational expressions should be captured in terms of empty elements in syntactic logical forms, *where syntax is conceived of as it standardly is*. Rather, the bindability of relational expressions should be captured in terms of empty elements in their own favored syntactic representations.

I do not myself believe that Discourse Representation Structure is the best account of natural language syntax; and I find Culicover and Jackendoff's talk of Conceptual Structure mysterious at best. None the less, I have, in this chapter, been as neutral as possible about what the correct syntactic theory of natural language is. All of these theorists account for relational expressions in terms of variables in their favored syntactic representations. Thus, besides a no doubt serious, though for these purposes irrelevant, disagreement about what counts as a suitable syntactic framework, there is, despite surface appearances, no dispute between these authors' conclusions and my own, construed in this second way.

## Conclusion

I have defended the thesis that all effects of extra-linguistic context are traceable to logical form. However, the considerations I have used are not always dependable. For example, it does not in general appear that possessive constructions, such as "John's book", involve a bindable variable whose values in different contexts are different salient relations. For example, "In most ways, John's book is nice" does not have a reading according to which the interpretation of the phrase "John's book" varies with the values introduced by the quantifier expression "most ways".<sup>43</sup> There are also other contextual phenomena, in particular focus, which need to be incorporated into a final account of these matters. But these are topics about which we are in any case in the dark. It would not do to rest an argument for the existence of unarticulated constituents on constructions the outlines of the ultimate analyses of which are unknown.<sup>44</sup> The argument for the existence of unarticulated constituents is only persuasive if it can be shown that it is methodologically implausible, for a range of different context-dependent constructions, to postulate variables in the logical form, the values of which are the desired contributions to truth-conditions. If so, then what we have seen is that no persuasive argument for the existence of unarticulated constituents has been provided.

Here is another consequence of the above discussion. Philosophers often turn to claims of hidden indexicality when faced with a philosophical quandary. For example,

<sup>43</sup> Surprisingly, possessives in post-copular position do seem to allow a bound reading, such as in "In most ways, that is John's book".

<sup>44</sup> Furthermore, it is not at all clear that the best syntactic and semantic theory for possessive relations will not postulate an empty element whose value is the possession relation. For example, in the best work on Possessives known to me, namely Barker (1995), possessive constructions involve an empty determiner in English, whose value in different contexts is a function of the possession relation salient in that context. So, on this account, possession relations are indeed traceable to an empty element in logical form.

according to Tyler Burge's well-known account of the strengthened liar paradox, the truth-predicate in fact is an indexical expression, whose extension varies from context to context. In the course of deriving instances of the strengthened liar paradox, context shifts in such a manner as to change the extension of the truth-predicate, and thereby vitiate the derivation. It is absolutely crucial to Burge's account that attributions of truth are due to the indexicality of "true", rather than the presence of a bindable variable in logical form, since otherwise the strengthened liar paradox could simply be reproduced.<sup>45</sup>

Similarly, according to one version of a contextualist response to skepticism, the word "know" is in fact an indexical expression, whose content varies from context to context. Relative to non-skeptical contexts, its content is a relation which holds between persons and those true propositions they believe, for which they have some minimal epistemic position. Relative to skeptical contexts, by contrast, it expresses a relation which holds only between persons and those true propositions they believe, with respect to which their epistemic position is very strong. According to leading proponents of contextualism, it is important to the doctrine that the epistemic standards are provided by the context of use, and not by the subject of the knowledge-ascription. It is therefore important to the doctrine that the word "know" is an indexical, rather than a non-indexical expression correlated with a variable in logical form that can be bound by a quantified expression in the subject position of a knowledge attribution.<sup>46</sup>

An additional consequence of the arguments I have given in the last section is to undermine the force of such appeals to "hidden" indexicality. If philosophically loaded expressions such as "true" and "knows" really were indexicals in the narrow sense of the term, then we should expect to find examples of such unobvious indexicality in the philosophically uncontroversial parts of our speech. However, what we have seen is that the vast number of cases of uncontroversial context-dependence do not involve indexicality, narrowly construed. Of course, obvious indexical expressions are indexicals, narrowly construed. But if words such as "true" and "knows" were hidden indexicals, then we should expect to discover cases of uncontroversial context-dependence that are best explicated in these terms. The fact that we do not provides some evidence that narrow indexicality is restricted to words such as "I", "here", and "now".

My central purpose in this chapter has been to explain and defend the thesis that all truth-conditional effects of context are traceable to logical form. If this thesis is

<sup>45</sup> "The indexical-schematic character of semantical predicates cannot be formally obviated by adding an argument place—relativizing them to a language, a level, a context, or a viewpoint. For quantification into the argument place will provide an open sentence just as subject to paradox as the 'naive' truth-predicate formalization" (Burge (1979, 192)).

<sup>46</sup> This is how I construe the emphasis on the importance of the "attributor" aspect of contextualism in DeRose (1999, sect. 4).

correct, then, after disambiguation, the process of interpreting a linguistic assertion has significant disanalogies with non-linguistic interpretation. Extra-linguistic context only can affect what is expressed in a linguistic assertion if its contribution can be traced to a constituent in the expression uttered. In contrast, the effects of extra-linguistic context on non-linguistic interpretation are constrained only by general considerations of relevance and rationality. Furthermore, we can maintain this strong distinction between linguistic and non-linguistic interpretation, without retreating from the thesis that semantic interpretation produces the full truth-conditions of utterances. Given their differences, we should therefore be suspicious of attempts to forge philosophically significant analogies between the different processes underlying the interpretation of linguistic and non-linguistic acts.<sup>47</sup>

## References

- AJDUKIEWICZ, K. (1967), 'Syntactic Connexion', in Storrs McCall (ed.), *Polish Logic* (Oxford: Clarendon Press), 207–31.
- AUSTIN, J. (1970), *How to Do Things With Words* (Cambridge: Harvard University Press)
- AUSTIN, J. (1979), 'Truth', in J. L. Austin, *Philosophical Papers* (Oxford: Oxford University Press), 117–33.
- BACH, K. (1982), 'Semantic Nonspecificity and Mixed Quantifiers', *Linguistics and Philosophy*, 4: 593–605.
- BACH, K. (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BACH, K. (ms) 'The Semantics Pragmatics Distinction: What it is and Why it matters'.
- BAR-HILLEL, Y. (1954), 'Indexical Expressions', *Mind*, 63: 359–79.
- BARKER, C. (1995), *Possessive Descriptions* (Stanford: CSLI)
- BARTON, E. (1990), *Nonsentential Constituents* (Amsterdam: John Benjamins).
- BLACKBURN, W. K. (1988), 'Wettstein on Definite Descriptions', *Philosophical Studies*, 53: 263–78.
- BURGE, T. (1979), 'Semantical Paradox', *The Journal of Philosophy*, 76: 169–98.
- CARSTON, R. (1991), 'Implicature, Explicature, and Truth-Theoretic Semantics', in Stephen Davis (ed.), *Pragmatics* (Oxford: Oxford University Press), 33–51.
- CASSAM, Q. (1997), *Self and World* (Oxford: Clarendon Press).
- CHIERCHIA, G. (1993), 'Questions with Quantifiers', *Natural Language Semantics*, 1: 181–234.
- CHIERCHIA, G. (1995), *Dynamics of Meaning* (Chicago: University of Chicago Press).

<sup>47</sup> I am grateful to Kent Bach, Lenny Clapp, Chris Collins, John Collins, Michael Glanzberg, Delia Graff, Richard Heck, Jim Higginbotham, Larry Horn, Jeff King, Peter Ludlow, Robert Stalnaker, Zoltán Gendler Szabó, Achille Varzi, and Timothy Williamson for comments and/or discussion. Various portions of this chapter were presented to the Philosophy Department at the University of Michigan and the Linguistics Department at Rutgers University. Comments by members of both audiences helped.

- CHOMSKY, N. (1995), *The Minimalist Program* (Cambridge: MIT Press).
- COOPER, R. (1996), 'The Role of Situations in Generalized Quantifiers', in Lappin (1996), 65–86.
- CRESSWELL, M. (1973), *Logics and Languages* (London: Methuen).
- CRESSWELL, M. (1996), *Semantic Indexicality* (Dordrecht: Kluwer).
- CRIMMINS, M. (1992), *Talk about Beliefs* (Cambridge: MIT Press).
- CULICOVER, P. AND JACKENDOFF, R. (1995), "'Something else" for the Binding Theory', *Linguistic Inquiry*, 26: 249–75.
- DEROSE, K. (1999), 'Contextualism: An Explanation and Defense', in John Greco and Ernest Sosa, *The Blackwell Guide to Epistemology* (Oxford: Blackwell), 187–205.
- ENGDAHL, E. (1986), *Constituent Questions* (Dordrecht: Reidel).
- EVANS, G. (1985), 'Pronouns, Quantifiers, Relative Clauses (I)', in Gareth Evans, *Collected Papers* (Oxford: Clarendon Press), 76–152.
- FIENGO, R., AND MAY, R. (1996), 'Anaphora and Identity', in Lappin (1996), 117–44.
- GRICE, PAUL (1989), 'Logic and Conversation', in *Studies in the Way of Words* (Cambridge: Harvard University Press), 22–40.
- HARMAN, G. (1972), 'Deep Structure as Logical Form', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (Dordrecht: Reidel).
- HEIM, I., AND KRATZER, A. (1998), *Semantics in Generative Grammar* (Oxford: Blackwell).
- HIGGINBOTHAM, J. (1986), 'Davidson's Program in Semantics', in E. LePore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (Oxford: Blackwell), 29–48.
- HIGGINBOTHAM, J. (1993), 'Interrogatives', in K. Hale and S. J. Keyser (eds.), *The View from Building 20: Essays in Honor of Sylvan Bromberger* (Cambridge: MIT Press), 195–227.
- HIGGINBOTHAM, J. (ms) 'A Plea for Implicit Anaphora'.
- HINTIKKA, J., AND SANDU, G. (1997), 'Game Theoretical Semantics', in van Benthem and Ter Meulen (1997), 361–410.
- JANSSEN, T. (1997), 'Compositionality', in van Benthem and Ter Meulen (1997), 417–73.
- KAPLAN, D. (1989), 'Demonstratives', in Almog, Perry, and Wettstein (1989), 481–564.
- KING, J. (1995), 'Structured Propositions and Complex Predicates', *Noûs*, 29: 516–35.
- KOOPMAN H. AND SPORTICHE D. (1982/3), 'Variables and the Bijection Principle', *The Linguistic Review*, 2: 139–60.
- LAPPIN, S. (ed.) (1996), *The Handbook of Contemporary Semantic Theory* (Oxford: Blackwell).
- LEWIS, D. (1983), 'General Semantics', in David Lewis, *Philosophical Papers, Volume 1* (Oxford: Oxford University Press), 189–232.
- MAY, R. (1985), *Logical Form: Its Structure and Derivation* (Cambridge: MIT Press).
- MITCHELL, J. (1986), 'The Formal Semantics of Point of View', unpublished UMass Amherst doctoral diss.
- PARTEE, B. (1989), 'Binding Implicit Variables in Quantified Contexts', *Proceedings of the Chicago Linguistics Society*, 25 (Chicago: University of Chicago Press), 342–65.
- PERRY, J. (1986), 'Thought without Representation', in *Supplementary Proceedings of the Aristotelian Society*, 60: 137–52.
- PERRY, J. (1997), 'Indexicals and Demonstratives', in Hale and Wright (eds.), *A Companion to the Philosophy of Language* (Oxford: Blackwell), 586–612.

- PERRY, J. (1998), 'Indexicals, Contexts and Unarticulated Constituents', in Aliseda, van Gabeek, and Westerståhl (eds.), *Computing Natural Language* (Stanford: CSLI), 1–11.
- RECANATI, F. (1993), *Direct Reference* (Oxford: Blackwell).
- RIZZI, L. (1986), 'Null Objects in Italian and the Theory of Pro', *Linguistic Inquiry*, 17: 501–57.
- RUSSELL, B. (1985), *The Philosophy of Logical Atomism* (LaSalle: Open Court).
- SALMON, N. (1986), *Frege's Puzzle* (Cambridge: MIT Press).
- SOAMES, S. (1987), 'Direct Reference, Propositional Attitudes, and Semantic Content', in *Philosophical Topics*, 14: 47–87.
- SOAMES, S. (1986), 'Incomplete Definite Descriptions', *Notre Dame Journal of Formal Logic*, 27: 349–75.
- SPERBER, D., AND WILSON, D. (1986), *Relevance* (Cambridge: Harvard University Press).
- STAINTON, R. (1994), 'Using Non-Sentences: An Application of Relevance Theory', *Pragmatics and Cognition*, 2: 269–84.
- STAINTON, R. (1995), 'Non-Sentential Assertions and Semantic Ellipsis', *Linguistics and Philosophy*, 18: 281–96.
- STAINTON, R. (1997), 'Utterance Meaning and Syntactic Ellipsis', in *Pragmatics & Cognition*, 5: 51–78.
- STAINTON, R. (1998), 'Quantifier Phrases, Meaningfulness "In Isolation", and Ellipsis', *Linguistics and Philosophy*, 21: 311–40.
- STALNAKER, R. (1970), 'Pragmatics', *Synthese*, 22: 272–89.
- SZABOLCSI, A. (1989), 'Bound Variables in Syntax (are there any?)', in Bartsch, van Benthem, and van Emde Boas (eds.), *Semantics and Contextual Expression* (Dordrecht: Foris), 295–318.
- VAN BENTHEM AND TER MEULEN (eds.) (1997), *Handbook of Logic & Language* (Cambridge: MIT Press).
- VAN DEEMTER, K., AND PETERS, S. (eds.) (1996), *Semantic Ambiguity and Underspecification* (Stanford: CSLI).
- VARZI, A. (1993), 'Do We Need Functional Abstraction?', in J. Czermak (ed.), *Proceedings of the 15th International Wittgenstein-Symposium, Hölder-Pickler-Tempsky*, Vienna, 407–15.
- VON FINTEL, K. (1994), 'Restrictions on Quantifier Domains', UMass diss.
- WILLIAMS, E. (1995), *Thematic Structure in Syntax*, Linguistic Inquiry Monograph, 23 (Cambridge: MIT Press).
- YANOFSKY, N. (1978), 'NP Utterances', *Papers from the Fourteenth Regional Meeting of the Chicago Linguistics Society* (Chicago: Chicago Linguistics Society), 491–502.

# 2

## On Quantifier Domain Restriction

(with Zoltán Gendler Szabó)

### I Introduction

In interpreting the utterances of others, we cannot rely exclusively on the permanent features of the words used. For it is often the case that the very same words could be used in a different context to communicate something different. *The problem of context dependence* is the problem of explaining how context contributes to interpretation, that is to the process of determining what a speaker meant by making a linguistic utterance on a certain occasion.

The topic of this chapter is *the problem of quantifier domain restriction*, which is a special case of the problem of context dependence. What is the problem of quantifier domain restriction? Consider the sentence:

(1) Every bottle is empty.

Suppose someone utters (1) in a conversation. It is unlikely that what she intends to convey is that every bottle in the universe is empty; she most likely intends to convey that every one of a restricted class of bottles (say, the bottles in the room where she is, the bottles purchased recently, etc.) is empty. And, if the context is right, she can succeed in communicating such a proposition. Permanent linguistic features of (1)—its phonological and morphological constituents, its syntactic structure, the meanings of the lexical items it contains—do not determine the proposition thereby communicated. They cannot do so, for these features are the same on every occasion when the sentence is used, but on most of those occasions the speaker would

Thanks to Michael Glanzberg, Kai von Fintel, Delia Graff, Richard Heck, Peter Ludlow, Marga Reimer, and Timothy Williamson for valuable comments and discussion. Our greatest debt is to Kent Bach, who provided numerous incisive written comments (not all of which we managed to incorporate) that significantly improved the paper.



communicate a different proposition by the sentence. The problem of quantifier domain restriction is a special case of the problem of context dependence, because, to solve it, we need to explain how context, together with permanent linguistic features of quantified sentences, helps determine the proposition conveyed by an utterance of such a sentence, a proposition in which the domains of the quantifier expressions are suitably restricted.

There are many accounts of the phenomenon of quantifier domain restriction. We provide here a survey of the space of possible analyses, together with a set of considerations designed to select the best from amongst them. We hope that our exhaustive discussion of this special case of the problem of context dependence may provide some guidelines for how to decide, for an arbitrary case of context-dependent discourse, whether it should be treated syntactically, semantically, or pragmatically. We have no general theory here. Rather, we have a set of distinctions and arguments which we hope will prove useful in future research into context-dependent phenomena.

The first two sections are devoted to articulating some distinctions relevant to the problem of context dependence generally. In the remainder of the chapter, we use these distinctions to address the problem of quantifier domain restriction.

## II The Foundational and Descriptive Problems of Context Dependence

Nothing short of an extremely comprehensive theory of linguistic communication could provide a general solution to the problem of context dependence. To explain successfully how, together with linguistic rules, context determines what people mean by the utterances they make requires a great deal of information about the mind and about language that is currently unavailable. But perhaps we can separate the intractably hard questions from the more or less manageable ones in this area. This is what we will attempt in this section.

Consider a simple and uncontroversial case of context dependence, when a speaker pointing at one of the tyres of his car utters the sentence:

- (2) That is flat.

The most obvious way in which the interpretation of this utterance depends on the context is that, in order to know what the speaker means, one has to know what is being demonstrated. It is uncontroversial that the full account of the context dependence of this utterance must include a specification of the truth-conditions of the utterance, and that this specification will, in turn, rely on a semantic clause like:

- (3) An occurrence of “that” in a context *c* refers to what is being demonstrated in *c*.

Context-dependent expressions, such as pronouns and demonstratives, behave in natural languages somewhat like variables in the language of predicate calculus. Variables in predicate calculus occur either bound or free. In the semantics of first-order logic, the interpretation of a bound variable is linked to the interpretation of a variable binding device; in standard first-order languages, a quantifier. Free variables do not receive an interpretation once and for all. Rather, their interpretation is relative to an ordered sequence of objects. In one standard treatment, variables receive indices; the variable marked with the index  $i$ , relative to a sequence, receives the  $i$ -th member of that sequence as value. In natural language, so-called “bound” pronouns are like bound variables insofar as their interpretation is linked to the interpretation of another expression—an antecedent or a quantifier. Other pronouns—the “unbound” ones—depend for their semantic value on the context in which they are used, and in this regard they are like the free variables of predicate calculus. Demonstratives and indexicals are parallel to unbound pronouns, and hence also to free variables.

Many theorists, impressed by the above analogy, have identified contexts with ordered sequences of objects, akin to the sequences in the semantics of first-order logic. Context-dependent expressions then take contexts as arguments, and yield the corresponding member of the context as value. The linguistic meaning of a context-dependent expression is represented, in such frameworks, as a function from contexts to values. This treatment of contexts has a distinguished recent history (cf. Montague (1974); Kaplan (1989a); Lewis (1970)). However, there are other traditions, motivated for the most part by problems distinct from that of supplying the values of demonstrative and indexical expressions. For example, to treat phenomena such as presupposition, it is helpful to view contexts as sets of propositions (cf. Stalnaker (1973, 1974)). One might find the representation of contexts as sequences of objects misleading, because one thinks that the work done by such entities is in fact done by sets of propositions. One might also find the treatment of contexts as sequences of objects objectionable on its own terms.<sup>1</sup>

We will not enter these debates about the proper representation of conversational contexts. That is a subject worthy, at least, of its own paper. Rather, we will take the notion of conversational context for granted, and speak of expressions having values

<sup>1</sup> According to Cresswell (1973, 111) identifying contexts with sequences of relevant objects is problematic, since it seems to require specifying in advance of interpretation a finite list of contextual coordinates which are relevant for the determination of content, and Cresswell suspects “that there is no such list” (1973, 111). Cresswell himself defines a context of utterance as a certain kind of property of an utterance. Kamp (1981a) objects to Cresswell’s approach, on the grounds that there is no principled way to distinguish the properties of utterances that, according to Cresswell, are contexts, from those that are not.

relative to contexts, or utterances communicating propositions relative to contexts, without giving a formal characterization of the notion (or notions) of context at issue.

There has been much debate in recent philosophy of language about whether one should consider the “true” bearer of propositional content to be a sentence relative to a context, or rather an *utterance* of a sentence. For example, according to David Kaplan (1989a, 522, 546), the former notion is required for model theoretic purposes. To evaluate an argument for validity, one needs to evaluate all of the premises and the conclusion with respect to the same context. But, Kaplan worries, on an approach that takes the true bearers of propositional content to be utterances, this sort of evaluation will be impossible, because distinct utterances occur in distinct contexts.<sup>2</sup>

This is also a debate we prefer to evade. For clarity’s sake, however, we fix on one vocabulary. We will formulate our semantic rules in terms of expressions relative to contexts. Furthermore, we follow Kaplan in speaking of the bearers of content as *occurrences* of expressions, by which we mean pairings of expressions and contexts. So, an occurrence of an expression *e* is a pairing of *e* and a context *c*. However, our usage of this vocabulary does not mean we have a commitment to one position or another in the debate about whether the ultimate bearers of content are expressions relative to contexts, or uses of expressions.

Given this background, we can divide (3) into two claims:

- (4) Relative to *c*, “that” denotes *o*.
- (5) *o* is *what is demonstrated* in context *c*.

These two claims correspond to two aspects of the problem of context dependence of the interpretation of an utterance of (2). Someone who accepts (4) as part of her account of this context dependence accepts (i) that the relevant context dependence is linked to the occurrence of the pronoun “that” within (2), and (ii) that the relevant context dependence is a matter of the semantic value of this pronoun. There is a sense in which these commitments jointly answer the question “how does the interpretation of an utterance of (2) depend on the context?” In another sense, however, a solution to the problem of context dependence requires more. For, only given (4), we do not know *what it is in virtue of which* *o* is the semantic value of “that” in *c*. (5) attempts to address this question. According to it, *o* is the semantic value of “that” in *c*, because it is the object demonstrated in the context in which the

<sup>2</sup> For more arguments in favor of an expression-based approach see Kaplan (1989b, 584–5) and Braun (1996, 152). Crimmins (1992, 195–6); Israel and Perry (1996); and Garcia-Carpintero (1998), have advanced considerations in support of a use-based approach.

utterance was made. Though (5) is no doubt far too simplistic, it addresses a question left unaddressed by (4).

Let us try to formulate in full generality the idea that led to the further analysis of (3). The *descriptive problem* of context dependence for an expression *e* relative to a context *c* is the problem of deriving the interpretation of *e* relative to *c*, given a prior characterization of what features of the context *c* have a bearing on the interpretation. (4) is a solution to the descriptive problem of context dependence for the demonstrative expression “that” relative to a context in which *o* is the value of that demonstrative expression. (4) is highly abstract, in the sense of not specifying what in the context makes it the case that the object mentioned in (4) is the denotation of the demonstrative in that context. The *foundational problem* of context dependence for an expression *e* relative to a context *c* is the problem of making these specifications. That is, the foundational problem of context dependence for an expression *e* relative to a context *c* is specifying what it is about the context in virtue of which certain entities (be they objects, properties, or propositions) play the role they do in the interpretation of an occurrence of *e*.

Separating the descriptive from the foundational problem of context dependence by distinguishing between (4) and (5) is customary in discussions of demonstratives. Semanticists often think of the semantic value of a demonstrative as “given” by the context, and relegate questions of what exactly it is in virtue of which it counts as the semantic value of that demonstrative in that context to a separate field of study.<sup>3,4</sup> One can see this distinction as a specific instance of a general division within semantics between descriptive and foundational work. Descriptive semantics gives the semantics for a language by assigning semantic values to simple expressions relative to contexts and by specifying rules that determine the semantic values of more complex expressions in terms of their simpler constituents. It does not attempt to decide what in the practice of the language users explains why that semantics is the right one for

<sup>3</sup> Cf. David Braun’s fairly typical attitude in the following remark: “I am imagining that in the formal semantics, the demonstratum of a context is simply given. But informally we can think of the demonstratum of a context as being determined, in some way that I won’t go into here, by one or more factors of the following kind: the speaker’s intentions, her overt and covert demonstrations, and the contextual cues that allow listeners to determine the referent” (1994, 209). Reimer (1998, 105–6) makes the same distinction with respect to quantifier expressions.

<sup>4</sup> One can also make the distinction between descriptive and foundational problems of context dependence with respect to theories which take contexts to be sets of propositions. For example, there are different positions, among those who think of contexts in this way, about what determines the set of propositions which is the context. Some hold that the context is the set of propositions which the speaker supposes are shared (e.g., Stalnaker (1973, 1974)). Others hold that the context is determined by the commonly shared assumptions of all of the conversational participants. Still others have argued that propositions can make it into the context set without being entertained by any of the conversational participants (cf. Gauker (1998)).

their language. Foundational semantics fills the explanatory gap left by descriptive semantics.<sup>5</sup>

In general, discussions of descriptive and foundational problems should be clearly distinguished from one another. To use an example discussed at length in Stalnaker (1997), in the case of proper names, an example of a foundational semantic debate concerns the viability of the causal theory of names, according to which the denotation of a name is due to the existence of a causal relation of the appropriate sort between it and its bearer. However, the causal chains are no part of the descriptive semantics of names. The descriptive semantics of names only involve linking them up with their bearers. Causal chains are rather part of a foundational semantic account of why names have the descriptive semantical interpretation they do. An example of a descriptive semantical debate is whether names are rigid designators, or are rather shorthand for non-rigid definite descriptions. Only confusion results from running such distinct questions together.

The distinction between foundational and descriptive semantics can be generalized to all aspects of the interpretation process. So, one can bracket foundational questions like: What exactly makes it the case that in the interpretation of a sentence *S* relative to a context *c*, we have to choose the wide scope, rather than the narrow scope reading? Or, why is it that under certain circumstances the use of the phrase “John’s bike” indicates that the relation between John and the bike is that of spatial proximity? Solutions to descriptive problems of context dependence should be thought of as ingredients in a descriptive theory; solutions to foundational problems belong to foundational accounts.

Most philosophical discussions of context involve foundational problems of context dependence. For example, debates about whether demonstrations or speaker intentions fix the reference of demonstrative expressions are instances of the foundational problem of context-dependence, as are debates about whether “I” must refer to the person who utters it.<sup>6</sup> Sperber and Wilson’s (1986) theory describes a general strategy exploited by language users to discover which features of the context are relevant for the resolution of ambiguity and semantic incompleteness, and so also involves the foundational problem of context dependence.<sup>7</sup> Gauker (1997) raises difficulties for possible solutions to the foundational problem of context dependence for quantifier domain restriction. And so on.

<sup>5</sup> Cf. Stalnaker (1997). See also Lewis’s (1975) distinction between a theory of languages and a theory of language, and Kaplan’s distinction between “semantics” and “metasemantics” (1989b, 573 ff.).

<sup>6</sup> For the former, see e.g. Kaplan (1978, 1989b); McGinn (1981); Wettstein (1984); Reimer (1991, 1992); and Bach (1992). For the latter, see for example Smith (1989).

<sup>7</sup> Indeed, Sperber and Wilson attempt to provide a very bold solution to the foundational problem of context dependence, since they argue that the same process underlies phenomena as distinct as the resolution of ambiguity and contextual supplementation of semantically incomplete information.

The reason that philosophers have focused their attention on foundational problems of context dependence is that they are sub-versions of the problem of linguistic intentionality, as are foundational semantic problems generally. Many philosophers who discuss foundational problems of context dependence seek to show that context-dependent constructions pose special problems for particular accounts of the problem of linguistic intentionality. For example, according to Griceans, the intentionality of language reduces to the intentionality of the mental. Philosophers have raised context-dependent constructions as counterexamples to this thesis.<sup>8</sup> According to these philosophers, the solutions to some foundational problems of context dependence involve non-mental facts, *contra* Gricean accounts of linguistic intentionality.

We will not engage in foundational debates about context dependence in this paper. Rather, our purpose is to resolve the often neglected descriptive problem of context dependence for quantifier expressions. However, the fact that we are addressing a descriptive problem of context dependence does not mean that we can completely neglect the corresponding foundational problem. Solutions to the descriptive and the foundational problems for a particular phenomenon of context dependence are interrelated, and so, even if one is primarily concerned about the descriptive problem, one should not devise a solution for it that would make any reasonable solution of the foundational problem impossible. In what follows, we will try to adhere to this general principle.

A full solution to a descriptive problem of context dependence involves two steps. As we will argue in the next section, there are three basic roles context plays in interpretation: grammatical, semantic, and pragmatic. The first step in resolving a descriptive problem of context dependence for a class of expressions is establishing which of these roles context fundamentally plays in the interpretation of occurrences of expressions in that class. The second step involved in a full solution to a descriptive problem of context dependence can only be undertaken once the first has been resolved. Once it has been decided whether context plays a grammatical, semantic, or pragmatic role in the interpretation of occurrences of these expressions, it must be decided how to integrate the relevant context-sensitivity into a grammatical, semantic, or pragmatic theory.

The first step in resolving the descriptive problem of context dependence for quantifier expressions requires deciding which role context plays in the provision of a restricted domain for a quantifier expression. We argue in sections IV–VI that quantifier domain restriction is a matter of semantics. Thus, the second step involves

<sup>8</sup> Usually, such philosophers argue their case using the examples of indexical and demonstrative reference, as in Wettstein (1984). However, Gauker attempts to formulate arguments for these conclusions with the examples of quantifier domain restriction (1997) and presupposition (1998).

deciding how to integrate the phenomenon into a semantic theory. This is the topic of our final section.

### III Three Roles of Context

In this section, we discuss three different ways in which interpretation depends upon features of context.<sup>9</sup> Interpretation, as we use this word, is the process of determining the information conveyed to an addressee by a speech act performed in a certain context. For the sake of simplicity, we will focus on the speech act of *assertion* and, among assertions, on what we will call *typical assertions*. In typical assertions (i) there is a single speaker and a single hearer, (ii) the speaker vocalizes a well-formed, meaningful sentence, and by doing so (iii) the speaker intends to convey a certain proposition. A typical assertion is successful just in case the hearer can identify the proposition the speaker intends to communicate. By focusing on typical assertions, we do not wish to downplay the theoretical significance of other speech acts. Instead, along with most linguists and philosophers of language, we hedge our bet on the question whether a theory of the interpretation of all speech acts can best be formulated on the basis of a theory of the interpretation of typical assertions.

Consider a situation in which a speaker performs a typical assertion. Suppose that the hearer perceives the voice of the speaker with perfect clarity and that the hearer has an ideal grasp of the phonology, morphology, syntax, and semantics of the language used by the speaker. There are, we believe, three distinctive ways in which in a situation like this communication may fail.

First of all, despite the clarity of the sound and despite his perfect linguistic knowledge, the hearer may not have enough information to identify the sentence uttered. For example, suppose the speaker produces certain sounds, much like what you would produce if you read the sentence (6) aloud:

(6) John likes to go to the bank.

<sup>9</sup> In making the distinctions which follow, we freely avail ourselves of some standard assumptions of linguistic theory. Here are two examples. First, we presuppose the viability of the standard picture of lexical and structural ambiguity, according to which ambiguous sentences correspond to distinct syntactic representations. This assumption is challenged in “underspecification” approaches, where, say, structural ambiguity is represented as a *semantic* phenomenon (e.g., Reyle (1993), and the essays in van Deemter and Peters (1996)); it is also questioned in May, 1985. A second example is our assumption in what follows that Verb Phrase ellipsis is a syntactic phenomenon, due to some sort of syntactic rule of reconstruction or copying, or PF deletion under a syntactic parallelism condition (Chomsky (1995, 125)). Most of the by now vast literature on ellipsis accords with this assumption. But the assumption certainly has had its detractors (e.g., Klein (1985); Dalrymple, Shieber, and Pereira (1991)).

There is a perfectly legitimate sense of “word” in which the hearer could complain that he does not know what the last word of the utterance was, and, consequently, does not know which sentence the speaker used. Here is another example. Suppose the speaker produces a sequence of sounds much like you would reading (7) aloud:

(7) Visiting friends can be annoying.

There is a legitimate sense of “grammatical structure” in which the hearer could say that he does not know which grammatical structure to ascribe to this sequence of sounds, and, consequently, does not know which sentence was uttered.

In the first of these examples, the reason the hearer may fail to know what sentence was uttered is that the sound that is produced by the speaker’s utterance of “bank” is linked to what any lexicographer would recognize as distinct lexical items. One lexical item means financial institution, and the other means the edge of a river.<sup>10</sup> In the second of these examples, the reason that the hearer may be ignorant of the sentence uttered is that the sound that is produced by an utterance of (7) can correspond to what any syntactician would recognize as two distinct analyses. One means that certain people—friends who come to visit—can be annoying; the other means that a certain activity—visiting one’s friends—can be annoying.

To clarify matters, we need to distinguish between two senses of the words that refer to linguistic expressions (words, phrases, sentences, etc.) Let us call an expression in the sense of a sequence of sounds a *phonological expression*, and let us call an expression in the special sense we have been motivating a *grammatical expression*. A phonological expression contains certain sounds in a certain linear order, but what exactly the relevant segments of the sound sequence correspond to in the lexicon, and how exactly their linear order translates to a grammatical structure plays no role in individuating the phonological expression itself. A grammatical expression, on the other hand, is either a lexical item or is constructed from lexical items arranged within a determinate syntactic structure. The sequences of printed letters displayed as (6) and (7) in this paper each correspond to a single phonological sentence but more than one grammatical sentence.

Having distinguished between phonological and grammatical expressions, we need to say a word about quotation. A string of letters enclosed between quotation marks may ordinarily designate either a phonological or a grammatical expression. We may

<sup>10</sup> It is common in lexicography to distinguish between two kinds of ambiguity: there are cases of polysemy, where the meanings associated with a single sequence of phonemes are related to one another, and cases of homonymy, where the meanings are independent. Polysemy is usually presented in dictionaries as sub-entries in a single entry; homonymy is presented as different entries. Etymology is a reliable, but by no means infallible guide in determining whether a certain sequence of phonemes is polysemous or homonymous. For example, there is an obvious etymological connection between the foot of a person and the foot of a mountain, but this is nevertheless arguably a case of homonymy.



with equal right talk about *the* sentence “Visiting friends can be annoying” and the sentences pronounced or spelled as “Visiting friends can be annoying”. To enhance clarity, in this chapter we use italics to refer to phonological expressions and quotation to refer to grammatical expressions. So, *bank* is a sequence of sounds and “bank” does not exist. What does exist is “bank<sub>1</sub>” and “bank<sub>2</sub>”, the former being a lexical item which means a financial institution and the latter another lexical item which means the edge of a river.<sup>11</sup> The use of other devices (such as indentation and numbering) will be left ambiguous.

In order to interpret typical assertions of others, we normally need to know what sentence they used, and know it in the stronger, grammatical sense of “sentence”. We will say that the phonological sentence produced by the speaker is *what is articulated* in the utterance, and that the grammatical sentence that in the context of the utterance is associated with the phonological sentence is *what is uttered* in the utterance. In cases like those discussed above, context has a characteristically *grammatical role* in determining what was uttered. Schematically:

what is articulated + context = what is uttered.<sup>12</sup>

The second way in which context can play a role in the interpretation of typical assertions concerns the identification of the proposition expressed by that which is uttered, relative to the context of use. Consider a case when someone utters (8):

(8) I am a philosophy professor.

Arguably, the sentence articulated carries no lexical or structural ambiguity, and hence corresponds to a single sentence uttered. Nevertheless, different occurrences of this single grammatical sentence may express different propositions. What is said by (8) relative to a context in which the speaker is Bill Clinton is different from what is said by (8) relative to a context in which the speaker is Robert Stalnaker, since the latter occurrence expresses a truth, while the former expresses a falsehood.

We use the expression “what is said” here in a somewhat arbitrary technical sense. It applies to propositions, entities which (like sentences) can be true or false and (unlike sentences) do not belong to any particular language.<sup>13</sup> We do not want to

<sup>11</sup> As will become evident below, our usage of quotation marks expresses Quinean quasi-quotation, rather than quotation proper.

<sup>12</sup> The intended interpretation of “ $x + y = z$ ” is that  $x$  and  $y$  jointly determine  $z$ . To say that what is articulated and context jointly determine what is uttered does not imply that in order to determine what was uttered one needs to know everything about the context. In fact, only a handful of contextual features are relevant. To determine exactly which ones is part of the foundational problem of context dependence.

<sup>13</sup> Some use the expression “what is said” to refer not just to propositions, but also to non-propositional entities, such as proposition radicals (e.g., Bach (1994)), which must be augmented with

commit ourselves here to anything more specific. In particular, we wish to remain neutral about the ontological status of propositions (are they *sui generis* or reducible to more fundamental entities?), about the structure of propositions (do they have constituents, and if they do how closely does their structure mirror the structure of sentences expressing them?) and about their role in the so-called propositional attitude ascriptions (are propositional attitudes two-place relations between a subject and a proposition or are there other argument places as well?). The only significant commitment we accept at this point is that propositions expressed by a typical assertion are individuated at least as finely as the truth-conditions of the utterance. That is, if two typical assertions differ in their truth-value potential, that is their truth-value in different possible situations, then they express different propositions.

In some semantic traditions, the first step in constructing a semantic theory for a language involves the provision of an algorithm that assigns propositions to sentences relative to contexts. In other semantic traditions, a semantic theory for a language involves giving a truth-definition for the language relative to contexts and perhaps also models. In the latter tradition, there might be no entity in the semantic theory that can be identified with a proposition. Which form a semantic theory ultimately should take is again a controversial question, which we do not wish to address in this chapter. Our discussion throughout is neutral as between these two approaches.<sup>14</sup>

Since what is uttered can be the same in different contexts, even though what is said differs, the former by itself cannot determine what is said. Rather, what is said by an occurrence of a grammatical sentence is dependent both on the semantic features of what is uttered *and* features of the context. Therefore, context does not just play a role in grammatical interpretation. There is also what we might call the *semantic role* of context, which is the role context plays in supplying additional values to what is said by an occurrence of a grammatical sentence, values that are not determined just from the linguistic meaning of the sentence uttered:

what is uttered + linguistic meaning + context = what is said.

Note that, according to our usage of the word “semantic”, any contextual provision to what is said counts as semantic. In this sense, semantics, by definition, is about propositions. However, there are other uses of “semantic”, with which our use should not be confused. For example, according to a second, equally legitimate use of the word, semantics is the study of the denotations of expressions relative to contexts. An object or property only enters into semantic interpretation by virtue of being either

extra constituents in order to be truth-evaluable. This usage must be sharply distinguished from our own.

<sup>14</sup> In particular, those who are not kindly disposed to proposition talk, such as Davidsonians, may construe our talk of propositions in terms of truth-conditions.

the semantic value of some expression relative to a context, part of the semantic value of some expression, or introduced via a rule of semantic composition. It is a substantial claim that these two uses of the term “semantic” cover the same territory. It might very well turn out that assigning denotations to the constituents of sentences relative to contexts and combining them often does not yield full propositional content, as has been argued by advocates of “the pragmatic view” (Travis (1997, 87)). However, this debate is irrelevant to our paper. As we use the term “semantic”, it includes any contextual contribution to the proposition expressed by an occurrence of a grammatical sentence, whether traceable to a constituent in what is uttered or not.

The grammatical role of context and the semantical role of context do not exhaust the interpretive roles of context. For even where a determinate proposition has been fixed that was expressed by the occurrence of a sentence, it may not yet be clear what proposition the speaker meant. For instance, someone may ironically utter:

(9) Fred is a fine friend

thereby communicating the proposition that Fred is a terrible friend. What this simple fact shows is just that, in the right context, words can convey something quite different from what they mean. Context plays what might be called a *pragmatic role*, in helping conversational participants move from what is said to what is communicated:

what is said + context = what is communicated.

Interpretation, in the broad sense in which we use this word, proceeds from the sentence articulated to the proposition communicated. Context plays a *grammatical role* in providing the proper lexical and syntactic analysis of the sentence articulated on a given occasion and thereby determining what was uttered. Context plays a *semantical role* in fixing what was said by that occurrence. Finally, context plays a *pragmatic role* in identifying the proposition communicated by the utterance. To solve the descriptive problem of context dependence for a particular expression, one needs to specify which of these three roles context plays in the interpretation of that expression relative to a context. This may not be more than a first step in solving the descriptive problem, but it is a necessary step.

One need not assume psychological reality for the grammatical, semantic and pragmatic *phases* of interpretation. Interpretation may not be a linear progression from the sentence articulated through the sentence uttered and the proposition expressed to the proposition communicated. But there are two assumptions that are immensely plausible. First, that in normal instances of *successful* communication, the hearer who grasps the proposition communicated will also know what sentence was uttered and what proposition was expressed by that sentence on the given occasion. Second, that in normal instances of *unsuccessful* communication the hearer might know the sentence uttered without knowing the proposition expressed or the

proposition communicated, or the hearer might know both the sentence uttered and the proposition expressed without knowing the proposition communicated. Cases when the hearer knows the proposition communicated without the proposition expressed or the proposition expressed without the sentence uttered are highly exceptional. The distinctions we drew in this section are based on nothing more than these two assumptions.<sup>15</sup>

#### IV The Problem of Quantifier Domain Restriction

The first step in resolving the problem of quantifier domain restriction is deciding whether the role played by context is syntactic, semantic, or pragmatic. Once this crucial issue is resolved, we can then turn to the question of which particular syntactic, semantic, or pragmatic account is correct. To set up our discussion of these issues, it would help to focus in on a representative example.

Suppose Lisa went to the store to buy some bottles to give to Max, who wanted to fill them with his home-made beer. Max asks whether the bottles Lisa bought need to be emptied first. In response, Lisa utters (1):

(1) Every bottle is empty.

In this situation, we can plausibly assume that by uttering (1) Lisa conveyed to Max the proposition that every bottle she just bought is empty. She succeeded in conveying this by relying, in part, on the context of her utterance. The first question is whether the role context played in the interpretation of Lisa's utterance was grammatical, semantic or pragmatic.

One way to think about the problem is this. Had Lisa been more explicit, she could have conveyed the same proposition by uttering (10) instead:

(10) Every bottle I just bought is empty.

The difference between the grammatical, the semantic, and the pragmatic solutions lies in the way they spell out the relationship between Lisa's actual utterance of (1) and her hypothetical utterance of (10). According to the grammatical approach, although the sentence articulated by Lisa is different from the sentence she would have articulated

<sup>15</sup> Many theorists would go further in their commitments. A fair number of successful natural language parsing programmes are designed assuming that interpretation has a grammatical, semantic and pragmatic phase and that completely separate computational routines are responsible for these phases (cf. Hirst (1987)). Many psycholinguists are committed to there being separate modules of the brain responsible for the grammatical, semantic and pragmatic aspects of interpretation. The jury is still out on whether such a modular approach is ultimately better than a non-modular one.

in the hypothetical case, there is no difference between what was and what would have been uttered; context supplies the additional material.<sup>16</sup> According to the semantic approach, the sentences uttered are different, but they express the same proposition. Finally, according to the pragmatic approach, the sentences uttered as well as the propositions expressed in the actual and the hypothetical situations are different. Nevertheless, what (1) and (10) communicate on these occasions is exactly the same.

In the abstract, there is not much that could be used to choose among these alternatives. We must therefore consider the merits and faults of specific ways of implementing the grammatical, semantic, and pragmatic approaches to quantifier domain restriction.

According to the grammatical approach to the problem of quantifier domain restriction, the way context operates to ensure that (1) and (10) express the same proposition in the envisaged scenario is by producing sentences whose context-independent features combine to express the same proposition. According to a plausible version of this approach, context simply provides an unarticulated portion of the sentence uttered. We call such accounts *syntactic ellipsis approaches*.

(11) is a standard example of syntactic ellipsis:

(11) Sam plays chess on Sundays. Max does too.

For any utterance of (11), it is plausible to say that although the words *plays chess on Sundays* are missing from what is articulated, “plays chess on Sundays” is nevertheless present in what is uttered. This claim can be supported by reflecting on what would be a *proper repetition* of this utterance. Suppose the addressee asks the question: “Excuse me, I did not hear you properly. What does Max do?” There is a sense in which (12) and (13) are appropriate answers to this question, whereas (14) and (15) are not.

(12) He plays chess on Sundays.

(13) Max plays chess on Sundays.

<sup>16</sup> One approach that does not meet this model, but in another sense might be thought of as a grammatical approach, is the *ambiguity* approach. According to this approach, the reason that (1) and (10) express the same proposition is that *bottle* is ambiguous. One of the lexical items corresponding to it is “bottle<sub>1</sub>” which is what we usually find as the first entry in any dictionary, and means what we would ordinarily expect. Normal dictionaries, however, omit the vast majority of other lexical items that correspond to the word *bottle*. One of the neglected lexical items is “bottle<sub>2067</sub>” which in the context of Lisa’s utterance means the same as the phrase “bottle<sub>1</sub> Lisa just bought”. When Lisa utters (1) in the situation described above, it is “bottle<sub>2067</sub>” that features in the sentence she uttered. If, however, she had uttered (10), “bottle<sub>1</sub>” would have been a constituent of the sentence she uttered. It is rather obvious why this approach has not much in its favour. We would have to assume that *bottle* corresponds to an infinite array of lexical items. If the meanings of these lexical items are unrelated, the lexicon is not learnable by a finite mind. If the meanings are related, it seems awkward to associate every one of them with a different lexical item. Approaches of this sort are generally implausible. In what follows, we only consider treatments which are not generally implausible.

- (14) He plays chess on those days of the week which follow a Saturday.  
 (15) Max plays some games of chess on Sundays.

One could take the analogous line with regard to Lisa's articulation of (1). The words "I just bought" are covertly present in the grammatical sentence uttered by her. The covert expression cannot be heard by anyone who listens to Lisa's utterance; it is a syntactic constituent that has no phonological manifestation. We will call this *the syntactic ellipsis theory of domain restriction*.<sup>17</sup>

We now turn to semantic approaches. What semantic approaches to the problem of quantifier domain restriction have in common is the commitment that although Lisa uttered different sentences in the actual and hypothetical case, the two sentences express the very same proposition in their respective contexts. For example, consider the sentence:

- (16) John is tall.

Without contextual background we only have a vague sense of what this sentence might say. In order to determine the proposition it expresses, we have to know what is the relevant comparison class with regard to which John is said to be tall. The very same sentence can express a truth if that class only includes John's colleagues and a falsehood if it also includes several professional basketball-players. There is, however, no reason to think that some natural language expression denoting this comparison class is present in the grammatical sentence "John is tall". Rather, one can say that the relevant class is the value of a *contextual parameter*, and that an occurrence of (16) expresses a proposition relative to this value.

There are two sorts of ways to incorporate this general account of comparative adjectives into systematic theory. According to the first, what is uttered contains a

<sup>17</sup> There are two versions of the syntactic ellipsis approach to quantifier domain restriction, which are usefully distinguished:

*Syntactic Ellipsis Theory of domain restriction, Version 1*

For any utterance of a phonological sentence S containing quantified expressions, the quantifier domain restrictions for the quantifier expressions in S are due to the presence, in the grammatical sentence uttered, of unarticulated natural language expressions whose interpretation is context-independent.

*Syntactic Ellipsis Theory of domain restriction, Version 2*

For any utterance of a phonological sentence S containing quantified expressions, the quantifier domain restrictions for the quantifier expressions in S are due to the presence, in the grammatical sentence uttered, of unarticulated natural language expressions.

According to the second, less restrictive version, unarticulated context-dependent natural language expressions may play a role in quantifier domain restriction. See the discussion of the "explicit approach" in Neale (1990, 95 ff.) for a discussion of these two versions of the syntactic ellipsis approach.

variable that is assigned a comparison class by context. For example, (16) might be informally represented as:

- (17) “John is tall F” is true relative to a context *c* if and only if John is tall for a member of the value of *F* relative to *c*.

where context would assign a comparison class to the variable *F*. According to the second, reference to the comparison class only appears in the metalanguage (cf. Heim and Kratzer (1998, 71)):

- (18) “John is tall” is true relative to a context *c* if and only if John is tall for a member of the comparison class provided by *c*.

On both of these accounts, an occurrence of (16) and an occurrence of:

- (19) John is tall for a basketball-player

may say the same thing, despite being distinct grammatical sentences. The role of context is not to provide expressions, but rather to provide semantic values.

One can follow a similar line in the case of quantifier domain restriction. There are no covert natural language expressions which context adds to what is articulated by Lisa’s utterance of (1), but there are covert semantic values which play their role in determining the proposition expressed. The semantic value of the sentence is a proposition that quantifies over the relevant bottles. This is due to the fact that the value of a contextual parameter somehow contributes to the semantic value of the whole sentence. We call any such approach a *semantic parameter approach* to the problem of quantifier domain restriction.

There are several ways to implement a semantic parameter approach, depending on where the contextual parameter is incorporated into the theory. As in the case of comparative adjectives, the main question is again whether the parameter is present in what is articulated, as in (20), or shows up only in the metalanguage, as in (21):

- (20) “Every bottle *F* is empty” is true relative to *c* if and only if every bottle in the domain *c* provides as the value of *F* is empty.  
 (21) “Every bottle is empty” is true relative to *c* if and only if every bottle in the domain provided by *c* is empty.

Again, in both cases, the role of context is not to supplement a phonological sentence with covert material to yield a distinct grammatical sentence. Context rather supplies semantic values.

One might feel that the approach in (20) is better called “grammatical”, since the quantifier domain is represented in the grammatical sentence itself. But this is not how we have chosen our terminology. According to grammatical approaches to the problem of quantifier domain restriction, Lisa’s actual utterance of (1) and her

hypothetical utterance of (10) correspond to the very same grammatical sentence. This is clearly not the case according to a theory that avails itself of a clause like (20). In such a theory, the grammatical sentence actually uttered by Lisa does not contain the lexical items “I”, “just”, and “bought”, but the grammatical sentence uttered by her in the hypothetical utterance does.<sup>18</sup>

Finally, there are ways to treat quantifier domain restriction pragmatically. According to a pragmatic theory of quantifier domain restriction, the actual occurrence of (1) and the hypothetical occurrence of (10) express different propositions. However, they *communicate* the same proposition. Here is an example of a case in which a pragmatic approach is plausible (cf. Bach (1994, 134)). Imagine a child who is crying because of a minor cut. Her mother attempts to calm her by uttering:

(22) You are not going to die.

Intuitively, the proposition communicated in this case is something like that expressed by an occurrence of “You are not going to die from that cut”. It is, however, possible that this is not what is expressed by (22) in this context. Perhaps what was said was that the child will not die *tout court*, which is quite certainly false.

There are ways to treat quantifier domain restriction in a similar fashion. For example, one may say that what is said by (1) in the envisaged context is the obviously false proposition that every bottle in the universe is empty. Realizing that the proposition is obviously false, her audience looks for a contextual elimination of the pragmatic anomaly. The reasoning to be followed here can be thought of as broadly Gricean. Upon realizing that the proposition expressed violates the maxim of quality, the interpreter looks for some other proposition that the speaker may reasonably be taken to have meant.<sup>19</sup> We call any such approach, a *pragmatic approach* to quantifier domain restriction.

As in the case of the expression “semantic”, our vocabulary here, too, could lead to misunderstanding. An approach does not count as pragmatic just because broadly Gricean mechanisms may be used to select a particular domain of quantification. For example, one could hold that broadly Gricean mechanisms are used to determine *what is said* by an occurrence of a quantified sentence. Such an approach counts as semantic, in our terms, not pragmatic. Rather, an approach is pragmatic just in case the quantifier domain restriction determined by contextual features does not affect what is said, but rather only what is communicated.

<sup>18</sup> A proposal according to which the ellided material contains *just* a covert indexical expression might be considered to be intermediate between a syntactic ellipsis and a purely semantic approach. However, we classify this proposal as a syntactic ellipsis proposal, and include discussion of it in our section discussing the syntactic ellipsis approach.

<sup>19</sup> However, the reasoning differs from standard cases of conversational implicature, since the proposition implicated in such cases will have to be semantically similar to the proposition expressed, whereas no such requirement exists in normal cases of implicature.



In the next two sections, we assess the merits of the different approaches to the problem of quantifier domain restriction. We argue first that the syntactic ellipsis approach is incorrect.<sup>20</sup> Next, we argue that all versions of the pragmatic approach are incorrect. It then follows that the role played by context in the provision of a contextually restricted domain to quantifier expressions is semantic.

## V Against the Syntactic Ellipsis Approach

According to the syntactic ellipsis theory, when someone is articulating a sentence like:

- (1) Every bottle is empty

what is uttered contains an expression which does not correspond to any segment of what was articulated. This unpronounced expression is a one-place predicate and its function is to restrict the domain of the quantifier “every” in (1).<sup>21</sup> So, the sentence uttered is “Every bottle which is *F* is empty”, where *F* abbreviates the unpronounced predicate. The domain of quantification is simply the intersection of the set of all bottles and the set of all things that are *F*. What expression *F* is depends on the context in which (1) is uttered.

What makes the syntactic ellipsis theory initially plausible is the idea that normal linguistic communication essentially relies on uses of linguistic expressions. Suppose a speaker utters the sentence “Snow is white” and thereby successfully conveys the proposition to a hearer that snow is white. The hearer came to know what the speaker meant because the hearer knows which words were used by the speaker, knows how one can properly use those words, and knows that the speaker used his words properly. In particular, the hearer came to know that the proposition the speaker meant is about snow because the speaker used the word “snow”, because the hearer knows what “snow” means, and because the hearer knows that the speaker meant just what the word means. So, it seems that a feature of the proposition conveyed (namely, that it is about snow) can be linked to certain features of the speaker’s use of a certain expression (namely, “snow”). Analogously, one might suggest that since an utterance of (1) can convey in the course of normal linguistic communication a proposition that is about a restricted class of bottles, there must be an expression in the

<sup>20</sup> Others who have criticized the syntactic ellipsis approach to quantifier domain restriction include Bach (1994, 130 ff.) and Recanati (1996, 448 ff.).

<sup>21</sup> This assumption is natural, but it is not indispensable. One could argue, for example, that the unarticulated expression is a relative clause, an adjectival phrase, or a prepositional phrase. None of these decisions makes a difference to our discussion in this section.

sentence uttered that is appropriately linked to that class. According to the proposal we are concerned with here, that expression would be “bottle which is *F*” for some contextually specified value of the schematic letter *F* and the appropriate linkage is that this predicate is true of all and only the elements of the class that the proposition conveyed is about.<sup>22</sup>

The main problem the syntactic ellipsis theory faces is that of *underdetermination*.<sup>23</sup> There are very few cases where there is a single plausible candidate for the role of the domain restricting predicate. Consider again the situation described in the previous section where (1) is uttered by Lisa. We assumed that the proposition conveyed by Lisa’s utterance was the same that she could have conveyed by uttering the sentence:

(10) Every bottle I just bought is empty.

It seems natural then to assume that *F* is simply the predicate “was just bought by me”. However, one might wonder whether this choice is somewhat arbitrary. Why could *F* not be the predicate “was recently purchased by me”, or “is one of those things that I bought at the store”? We need not even restrict ourselves to synonyms or near synonyms. It seems that, in the context of Lisa’s utterance, *F* could equally be something like “is one of those things you are looking for”. Since the sole function of *F* in communicating is to restrict the domain of the quantifier, it is hard to see how to select among predicates that apply to the same bottles.

We can articulate the problem of underdetermination using the distinction between the descriptive and the foundational problems of context dependence. The concern is that the syntactic ellipsis approach to quantifier domain restriction provides a solution to the descriptive problem by placing intolerable burdens on any possible

<sup>22</sup> According to the “modified traditional account” of Blackburn (1988), an utterance of (1) expresses the set of propositions expressible by sentences which result from (1) by expanding the quantifier phrase “every bottle” into an expression that the speaker “would be prepared to fall back on” (Blackburn (1988, 271)). We are not sure what is meant by this latter phrase. Be that as it may, Blackburn’s account is not a version of a syntactic ellipsis approach. In cases of syntactic ellipsis, there is a unique phrase recoverable from the context.

<sup>23</sup> A contemporary paper in which an underdetermination problem is raised is in Wettstein (1981), in which it is used to argue that definite descriptions cannot be treated along Russellian lines. We object to this usage of the problem of underdetermination on two grounds. First, it presupposes that the only strategy available to the Russellian to address the problem of quantifier domain restriction is the syntactic ellipsis approach. Secondly, the problems with the syntactic ellipsis approach are not unique to the definite description, but occur with all quantificational expressions. Schiffer (e.g. 1992, 512–18; 1995) has emphasized that underdetermination worries are relevant wherever “hidden indexicality” is at issue. His term for the problem of underdetermination is “the meaning-intention problem”. We have learned much from Schiffer’s discussions. However, his vocabulary presupposes the Gricean thesis that the intentions of conversational participants are all that is relevant to resolving the context dependence. Since we wish to remain neutral on issues concerning the foundational problem of context dependence, we have replaced his vocabulary by our own.

solution to the foundational problem. If context has to provide a specific predicate whose extension will contribute to the determination of the domain, a solution to the foundational problem involves specifying the relevant features of the context which select the predicate *F* among other candidates. And it is exceedingly hard to see what feature of the context could do that.

There is only one plausible principle that could tell us which predicate to choose in interpreting utterances of quantified sentences. The principle is that one should choose a purely demonstrative predicate, for there is a clear sense in which that is the simplest of all the acceptable contenders. So, the sentence uttered by Lisa would have to be “every bottle which is one of those is empty”. The demonstrative pronoun “those” will denote, relative to the context of Lisa’s utterance, the set of things that comprise the domain of quantification. This is the *pure indexical* version of the syntactic ellipsis theory.

The first thing to say about this proposal is that while it is syntactic in letter, it is semantic in spirit. Although there is a grammatical expression to be identified on contextual grounds, the expression is almost completely void of context-independent meaning. What ultimately settles the domain of the quantifier is the semantic mechanism that assigns the appropriate value to this expression. However, a defender of the pure indexical version of the syntactic ellipsis theory may concede that the proposal is very close to a semantic theory but insist that it is superior to those because it links the domain of the quantifier to a particular English expression within the sentence.

But there is still perhaps an argument to be given against the pure indexical approach. Suppose that Max is not a fully competent speaker of English. In fact, he started to learn the language only a few weeks ago. As it happens, the first lessons in his language book focus on prepositions and the use of “every” and “some” as well as a few basic nouns like “bottle”. The use of demonstrative pronouns is not discussed until unit 7 and Max is not there yet. We believe that under these circumstances Max could grasp the proposition meant by Lisa *in the normal way*. This, of course, is in conflict with the pure indexical version of the syntactic ellipsis theory. Since Max does not know the word “those”, he cannot identify the sentence uttered by Lisa which contains that word as an unarticulated constituent. Consequently, he cannot in the normal way know what was expressed by Lisa’s utterance, and *a fortiori* he cannot in the normal way know what she meant.

This concludes our arguments against the syntactic ellipsis theories of quantifier domain restriction. The basic trouble with these theories can be summarized as follows. Either there is a way to specify which is the unarticulated constituent of the sentence uttered or there is not. If there is not, the hearer cannot know what sentence was uttered and consequently cannot know in the normal way what proposition was meant. That is, when a quantified sentence is used whose domain is not articulated,

the hearer can never know in the normal way what is conveyed, which is absurd. If there is a way to specify the unarticulated constituent, then it is conceivable that the hearer does not know this expression but knows all the articulated constituents of the sentence uttered. In such a case the syntactic ellipsis theory entails that the hearer cannot know in the normal way the proposition conveyed, which is implausible.

## VI Against Pragmatic Approaches

According to pragmatic approaches to quantifier domain restriction

- (1) Every bottle is empty

expresses, in every context, the false proposition that every bottle in the universe is empty. The audience then uses general pragmatic principles to infer from information available in the context that the speaker intended to communicate a proposition concerning a more restricted domain of quantification. Thus, occurrences of (1) always express false propositions, but they communicate true ones.

This characterization of the pragmatic approach may be somewhat narrow on two grounds. First, one might propose a pragmatic solution to the descriptive problem of quantifier domain restriction without claiming that utterances of sentences containing quantifiers are usually false. There might be some other peculiarity about them that would explain why the proposition expressed by these utterances differs from the proposition they convey. This sort of solution may be plausible for other instances of the problem of context dependence. It is, however, in the particular case of quantifier domain restriction, unclear to us what this alternative peculiarity could be.

Secondly, one might suggest that the pragmatic rules that help one to derive the proposition conveyed are not general principles of conversation. Domain restriction may well be something like conventional, rather than conversational, implicature. Note, however, that standard examples of conventional implicature (e.g. the implicature that when two clauses are conjoined with “but” there must be some contextually salient contrast between their contents) tend to add extra information to the proposition expressed, rather than override what is said. This would not be the case for a parallel account of quantifier domain restriction. So, the defender of such a conventional implicature approach has to provide an explanation for the distinction. Be that as it may, the assumption that a pragmatic approach would use general principles of conversation in deriving the proposition conveyed for utterances containing quantifiers is not an assumption that is essential to our arguments in this section.

Pragmatic approaches have an obvious advantage and an obvious disadvantage. The obvious advantage is that one can propose a syntax and semantics for sentences containing quantifiers that is extremely simple and does not involve covert expressions or covert semantic values. In this regard they follow the advice of radical pragmatics: try to keep your syntax and semantics as simple as possible. The obvious disadvantage is that one has to abandon ordinary intuitions concerning the truth or falsity of most sentences containing quantifiers. This is worrisome because accounting for our ordinary judgements about the truth-conditions of various sentences is the central aim of semantics. Since these judgements are the data of semantic theorizing, we should be careful with proposals that suggest a radical revision of these judgements. How to weigh the obvious advantage and the obvious disadvantage is not entirely clear. So, we will focus on arguments for and against pragmatic approaches that are independent of these considerations.<sup>24</sup>

Kent Bach is the philosopher most associated with pragmatic approaches to quantifier domain restriction. But it should be noted that Bach, especially in his most recent work, is not a straightforward proponent of pragmatic approaches as we have defined them. For Bach's distinction between semantics and pragmatics differs sharply from our own (cf. Bach (1999)). Bach uses the term "semantic" in the second sense discussed in section 3. Because of his belief that the effects of context on the truth-conditions of utterances usually outstrip denotation assignment and composition, Bach believes that the output of semantic interpretation is often a non-propositional entity. In addition to its standard Gricean role, for Bach, pragmatics involves turning these non-propositional entities into propositions. However, pragmatic approaches to quantifier domain restriction, in our sense, are ones that account for quantifier domain restriction in terms of Gricean-like inferences from the proposition expressed to the proposition communicated. It follows from this difference that Bach has resources at his disposal that a proponent of pragmatic approaches, in our sense, does not.

Bach is nevertheless an appropriate target for our arguments. First, in the special case of quantifier domain restriction, he in fact holds that semantics does deliver a proposition, that is, a truth-evaluable entity, and pragmatics yields a different one. Secondly, our positive argument against pragmatic approaches in our sense ultimately shows that quantifier domain restriction is due to denotation assignment to a quantifier domain variable in the actual syntactic structure of quantified sentences, and so is semantic in Bach's sense.

Here is an argument, due to Bach (1994, 138–9), against the view that one should incorporate quantifier domain restriction into the syntax or the semantics of

<sup>24</sup> Thus, our argument against pragmatic approaches will differ substantially from that given by Recanati, in Recanati (1993). Recanati's critique of pragmatic approaches focuses on essentially this consequence of them (cf. his discussion of the "Availability Principle", 248 ff.).

quantification. The first premise is what Bach calls “syntactic parallelism”. According to it, syntactically parallel sentences should receive parallel semantic treatment. Bach’s second premise is that there is no reason to accept that the proposition expressed by utterances of (23) quantifies over a restricted domain:

(23) A book is on the table.<sup>25</sup>

Bach’s third premise is that (23) is syntactically parallel to any sentence which results from replacing “a book” by any other quantifier expression. Bach’s conclusion is then that there is no reason to believe that the proposition expressed by a use of (1) quantifies over a restricted domain.

The first problem with this argument is that syntactic parallelism cuts both ways. Bach’s opponent believes that utterances of (1) are best interpreted as quantifying over a restricted domain. So, if the opponent accepts Bach’s first and third premises she can justifiably say that the second premise cannot be true. There *is* a reason for thinking that utterances of (23) express a proposition that quantifies over a restricted domain, namely, that utterances of (1) do so, and sentences (1) and (23) are syntactically parallel.

The second problem with Bach’s argument involves its second premise. What Bach needs to show is that there are reasons to believe this premise that could be persuasive to someone who does not already believe the conclusion. What could such a reason be? We suspect it is based on the observation that whenever an existentially quantified sentence is true when interpreted as governed by a domain, it is also true when interpreted as governed by a broader domain. Hence, it is true when interpreted as governed by a domain that encompasses absolutely everything.<sup>26</sup> This is a feature of existential quantification that is not shared by universal quantification: if a universally quantified sentence is true when interpreted as governed by a domain, it may nevertheless be false when interpreted as governed by a broader domain. So, it may seem to follow that assuming a more restricted domain for the proposition expressed serves no theoretical purpose.

But this line of reasoning is incorrect. It is clear that as long as we are convinced that an occurrence of (23) expresses a truth, we might as well interpret it as quantifying over absolutely everything. But what if, intuitively, an occurrence of (23) expresses something false? If (23) expresses a falsehood when interpreted as quantifying over  $D$ , it may still express a truth when interpreted as quantifying over a superset of  $D$  (that

<sup>25</sup> We ignore, in what follows, the context dependency associated with the quantifier expression “the table”.

<sup>26</sup> This feature is usually called “upward monotonicity”. A unary quantifier  $Q$  is upward monotone (or monotone increasing) iff for all sentences  $S$  whose main connective is  $Q$  and for all domains  $D$  if  $S$  is true with regard to  $D$  and  $D \subseteq D'$  then  $S$  is true with regard to  $D'$ .

is, the existential quantifier is not downward monotone). So, domain restriction is truth-conditionally relevant for existentially quantified sentences.

Consider the following example. John and Bill are printing copies of *Naming and Necessity* in their printing shop. There are thousands of copies of this book lying around. Lunch break is approaching and John complains to Bill that he wants to read a book, since he needs to get his mind off *Naming and Necessity*. Bill believes that there are several detective novels lying on the table beyond him, and, on this basis, utters (23). If, however, all there are on the table behind Bill are more stacks of *Naming and Necessity*, then this occurrence of (23) seems false. Intuitively, that is because (23), relative to this context quantifies over (copies of) books other than *Naming and Necessity*. Extending the domain in such a way that it would include copies of *Naming and Necessity* would result in a different interpretation for (23) and under this interpretation it would express a truth.

Now, Bach can (and would) deny that (23) expresses a falsehood in the context of John's utterance. But notice that the burden of proof is on him. In order for his argument to work, it must be the case that, intuitively, there are no truth-conditionally relevant effects of domain restriction on (23). He must believe that this alleged contrast between (1) and (23) can be appreciated independently of accepting the conclusion of Bach's argument, namely that quantifier domain restriction is pragmatic. But this is not true. There is no reason for Bach's opponent to accept the second premise of his argument. Therefore, Bach's argument is powerless against its intended audience.

It is one thing to undermine an argument in favour of pragmatic approaches; it is another to provide an argument against them. We now turn to the latter task.<sup>27</sup> What we claim is that the phenomenon of *quantified contexts* poses an insurmountable difficulty for pragmatic approaches. Quantified contexts are cases involving sentences containing multiple quantified expressions whose intuitive readings are only possible to capture by assuming that an index representing the quantifier domain of the second quantifier expression is bound by the first quantifier expression. Since the pragmatic approach does not postulate syntactically represented, or semantically reflected quantifier domains, it cannot capture these readings.

Consider the following sentences (see von Stechow (1994); Cooper (1993)):

- (24) In most of John's classes, he fails exactly three Frenchmen.
- (25) In every room in John's house, every bottle is in the corner.
- (26) Whatever John does, most of the class falls asleep.

In each of these examples, the domain of the second quantifier expression varies with the values of the first quantifier expression. For example, the proposition intuitively

<sup>27</sup> The arguments to follow are developed in greater detail in Chapter 1.

expressed by an utterance of (24) is the proposition that, for most  $x$  such that  $x$  is a class of John's, John failed three Frenchmen in  $x$ . Thus, the domain of the quantifier expression "three Frenchmen" varies with the value of the variable introduced by the quantifier "most". Thus, the quantifier domain variable associated with "three Frenchmen" is *bound* by the preceding quantifier expression. Similarly, the proposition intuitively expressed by an utterance of (25) is the proposition that in every room  $x$  in John's house, every bottle in  $x$  is in the corner. Therefore, the quantifier domain variable associated with "every bottle" is bound by the preceding quantifier expression. Finally, the proposition intuitively expressed by an utterance of (26) is the proposition that whatever action  $x$  John undertakes, most of the class in the situation in which  $x$  occurs falls asleep. Therefore, the quantifier domain variable associated with "most of the class" varies with the value of the variable introduced by the quantifier expression "whatever John does". Thus, the quantifier domain variable associated with "most of the class" is *bound* by the preceding quantifier "whatever".

However, the pragmatic approach does not posit any quantifier domain variable associated with the quantifier "three Frenchmen". According to the pragmatic approach, the only reading of (24) is one on which the second part of the sentence is completely unrelated to the first part of the sentence. Indeed, it is not clear, on a pragmatic approach, that sentences such as (24)–(26) express coherent propositions at all.<sup>28</sup>

Since it is not clear that (24)–(26) even express coherent propositions on the pragmatic approach, one standard defence of the pragmatic approach fails. According to this defence, occurrences of (1) *really* express a proposition about every bottle in the universe, since it is always possible to respond to an utterance of (1) by saying "Strictly speaking, that is not true, since there are some bottles in the universe that are not empty". According to the defender of the pragmatic view, what this shows is that, strictly speaking, occurrences of (1) express the proposition that every bottle in the universe is empty.

This defence is unpersuasive even for examples such as (1). For the defender of grammatical or semantic approaches to quantifier domain restriction may simply maintain that, in so responding to an utterance of (1), one thereby shifts

<sup>28</sup> In the face of this difficulty, a defender of the pragmatic approach may argue that although "he fails exactly three Frenchmen" does not contain any syntactically represented quantifier domain, this is not a problem since the full structure of (24) is akin to that of "For most  $x$  such that  $x$  is a class of John, he fails three Frenchmen in  $x$ ". There are two responses. The first is that it is unclear how this approach would treat examples such as (26), since it is completely unclear what the unarticulated constituent would be. Secondly, this sort of move is blatantly inconsistent with the spirit of the pragmatic approach. It is in part because the proponent of the pragmatic approach believes that it is absolutely illegitimate to postulate structure on semantic grounds that she is able to make the case for a pragmatic approach (cf. Bach (1994, 130 ff.)).



the context.<sup>29</sup> Be that as it may, even assuming that this line of defence is plausible for examples like (1), it is simply unavailable for examples of quantified contexts. For it is not coherent to respond to an utterance of (24) by saying, “strictly speaking, that is not true, since John has failed many more than three Frenchmen”.

Here is a possible response to our argument for the existence of quantifier domain variables. One might respond by conceding that in sentences such as (24)–(26) there is a quantifier domain variable that is bound by the initial quantifier, but deny that in an “unembedded” sentence such as:

(27) John failed exactly three students

there is a quantifier domain variable present. According to this response, a variable is associated with quantifier expressions only in the special case of bound readings such as (24)–(26).

However, consideration of facts from ellipsis serves to dispose of this response. Consider the discourse:

(28) John failed exactly three Frenchmen. In fact, in most classes John has taught, he has.

The natural reading of the second sentence in (28) is that in most classes  $x$  such that John has taught  $x$ , he has failed exactly three Frenchmen in  $x$ . However, if there is no quantifier domain variable present in the initial sentence in (28), then there is no way of deriving the natural reading of the second sentence.

The second sentence of (28) is a standard case of syntactic ellipsis (verb phrase ellipsis). According to standard theories of such ellipsis, the predicate “failed exactly three Frenchmen” in the first sentence is copied or reconstructed in the final syntactic structure of the second sentence.<sup>30</sup> If there is no quantifier domain variable available for binding in the predicate “failed exactly three Frenchmen” in the first sentence of (28), then the result of copying or reconstructing it in the logical form of the second sentence will also not contain a bindable variable, in which case there will be no way to derive its natural reading (see pp. 60–1). Therefore, on the assumption

<sup>29</sup> Indeed, the smoothest account of operators such as “strictly speaking” suggests that they are context-shifting *operators*. David Kaplan has, however, argued that there are no operators which shift contexts (Kaplan (1989a, 510 ff.)); see also, for discussion of Kaplan’s thesis, Israel and Perry (1996). We do not know what to think of Kaplan’s restriction on such operators.

<sup>30</sup> The argument to follow does not depend upon a copy theory of ellipsis; it would work equally well under the minimalist assumption that ellipsis amounts to PF deletion under a parallelism requirement.

that standard theories of syntactic ellipsis are correct, there is a bindable variable for quantifier domains present even in sentences such as (27).<sup>31</sup>

We have therefore demonstrated the existence of quantifier domain variables, or at least some process that semantically mimics syntactic binding.<sup>32</sup> In the face of these considerations, arguing that quantifier restriction must be treated pragmatically is akin to arguing that the reading in which “his” is bound in:

(29) Every boy loves his mother

should also be captured pragmatically. Examples such as (24)–(26) also pose an even greater difficulty for the task of developing a pragmatic approach to the problem of quantifier domain restriction. For it is even more difficult to see how to derive the intended meaning from the absurd proposition predicted to be expressed on such an approach.

We have argued that both the syntactic ellipsis approach and the pragmatic approach to quantifier domain restriction are incorrect. Thus, quantifier domain restriction must be, in our terminology, a semantic process. In the next section, we turn to the different candidate semantic processes that could underlie the phenomenon of quantifier domain restriction.

## VII Semantic Approaches

In this section, we discuss ways to represent quantifier domain restriction semantically. Our desire throughout is to keep our presentation both as informal and as theory neutral as possible. But we will of course need some rudimentary semantic tools. The phenomenon of quantifier domain restriction has received interesting treatment in

<sup>31</sup> Kent Bach has suggested to us another line of response, one which has some initial promise for examples such as (24) and (25), but is unavailable for examples such as (26). According to this suggestion, (24) is transformationally derived from:

(24') He fails exactly three students in most of John's classes.

However, (24) is not transformationally derived from (24'). (24') does not permit co-indexing between “he” and “John”; the theoretical explanation of this fact in syntax is that such co-indexing would violate Condition C of the Binding Theory. If the prepositional phrase “in most of John's classes” were to move in the way suggested, then one should not expect a reading of (24) in which “John” and “he” are co-indexed, because that would result in a strong crossover violation. Since (24) has a very natural reading in which “John” and “he” are co-indexed, (24) is not transformationally derived from (24').

<sup>32</sup> For the details of the latter sort of process, see Max Cresswell's treatment of examples such as (25) in Cresswell (1996, 81–7).

several different theoretical frameworks.<sup>33</sup> However, to meet our goals of maximum accessibility and theoretical neutrality, we will use only some elementary resources of model theory. Where possible, we provide explanations of the basic tools we use.

We also need, at the outset, to introduce some very rudimentary syntax. The first reason is that, to differentiate between proposals, we will need to talk about the structure of the sentences that are being interpreted by the semantics. But we also believe, and assume in this discussion, that semantic interpretation is run off the output of a syntactic mechanism, whose nature it is the purpose of syntactic theory to uncover.

We realize that some semanticists and philosophers use their semantic theories to interpret structures that differ greatly from the syntactic structures produced by plausible syntactic theories for natural language. There are essentially two different theoretical motivations for doing this: one that is popular among philosophers doing semantics and another that is popular among linguists doing semantics. The philosophical motivation is the felt need for regimented discourse. As far as pure syntactic evidence is concerned, it is hard to see what difference there would be between sentences like “2 is an even number” and “Joe is a skilful painter”. But then if semantic interpretation uses this shared syntactic structure as its input, it is hard to see how to avoid the conclusion that just as the truth of the second sentence commits us to the existence of Joe, the truth of the first sentence commits us to the existence of 2. And for many philosophers, this seems unacceptable.

In general, it is hard for philosophers to part from the idea that one can freely construct alternative semantic structures for various natural language sentences without being constrained by empirical evidence from linguistics. Such a view, however, is tantamount to the endorsement of the hypothesis that syntax is a superficial feature of language, detached from the way we understand the utterances of others. We find this hypothesis implausible in the extreme.

The linguistic motivation for being cavalier about syntax is the belief that semantic interpretation should take cross-sentential phenomena into account, and hence cannot run off the output of a syntactic mechanism that consists of isolated sentential structures. Instead, a proper semantic theory should proceed in two steps: first, articulating a mechanism that builds unified semantic representations out of syntactic structures associated with each of the sentences within a longer discourse, and second, assigning semantic interpretations to these structures.<sup>34</sup> Now, one can of course provide unified syntactic representations for sequences of sentences by providing the

<sup>33</sup> For example, Robin Cooper (1993) treats quantifier domain restriction in the “Extended Kamp Notation” of Barwise and Cooper (1993). Max Cresswell treats quantifier domain restriction in a “propositional language”, which mirrors variable binding with the use of sentence operators (cf. Part II of Cresswell (1996)).

<sup>34</sup> This is essentially the programme of Discourse Representation Theory (cf. Kamp (1981*b*); Kamp and Reyle (1993)).

standard syntactic representation for each sentence, and grouping together sequences of such representations under a text or discourse node (cf. Heim (1982)). Why not take these as the inputs of the assignment of semantic values? The reason might be that these syntactic representations are not interpretable compositionally, or that the resulting interpretation is too complicated, or perhaps that it lacks the psychological reality the unified semantic representations have.

Whether any of these claims is in fact true is an open empirical question.<sup>35</sup> However, the default assumption should certainly be that they are false. A system in which the outputs of the standard syntactic process are directly assigned semantic interpretations is significantly less complex than one which includes, in addition, a translation of these outputs into intermediate semantic representations, which are then interpreted by the semantics. Unless this additional complexity is justified by strong empirical arguments, good methodology should lead us to reject it.<sup>36</sup> That is our justification for our assumption that what is semantically interpreted are the outputs of a syntactic mechanism, whose nature it is the purpose of syntactic theory to describe.

We will call the output of the syntactic process that is visible to semantic interpretation a logical form. A logical form is a lexically and structurally disambiguated ordered sequence of word types, where word types are individuated both by semantic and syntactic properties. Logical Forms are phrase-markers. An example of such a phrase marker, for the sentence “Hannah loves Sue” is as follows:




Figure 1

The nodes in this diagram are the points labeled either with syntactic categories or lexical items. So, “Hannah” labels a node, as does “N”, “VP”, and “V” (we will also talk of labels of nodes as *occupying* these nodes). The nodes are connected by branches,

<sup>35</sup> The first of these claims was, in fact, the contention of Kamp (1981*b*). In the ensuing years, a number of authors have elaborated frameworks within which the phenomena which motivated the development of DRT receive compositional treatments (e.g. King (1987); Groenendijk and Stokhof (1991); Neale (1990)).

<sup>36</sup> We do not wish to be taken as denying the importance of approaches such as DRT. Being more flexible than conventional theories, they can approach phenomena which are as yet unapproachable by other means. That is, they articulate theoretical challenges for what might turn out to be more methodologically sound syntactic and semantic theorizing.

which are the lines in the diagram. We say that a node  $X$  *dominates* another node  $Y$  in a phrase marker if there is a path of branches leading downward in the tree from  $X$  to  $Y$ . We say that a node  $X$  immediately dominates another node  $Y$  in a phrase marker just in case  $X$  dominates  $Y$ , and there is no node  $Z$  between  $X$  and  $Y$ . Nodes that dominate other nodes are called *nonterminal nodes*. Nodes that dominate no other nodes are *terminal nodes*. The nodes labeled with lexical items such as “Hannah” or “loves” are always terminal nodes. These are the objects that we assume are interpreted by semantic theories.

The main division between semantic parameter approaches lies in whether the contextual parameter is provided as the value of a variable in the logical form of a sentence relative to a context, or whether it is provided in the metalanguage. Approaches of the latter sort may be called metalinguistic semantic parameter approaches. We begin by presenting and rejecting the two conceivable metalinguistic parameter approaches to quantifier domain restriction.<sup>37</sup> We then turn to different versions of the first sort of semantic parameter approach, and select among them.

Perhaps the most natural way of representing domain restriction semantically is in fact a metalinguistic semantic parameter approach. It is as follows. In a model-theoretic semantics, the truth of sentences is considered relative to a model. Each model has a domain of individuals, and an assignment of subsets of the domain to predicates. This suggests the following elegant treatment of quantifier domain restriction. Since the truth of a sentence is considered with respect to a model, and the model already has a domain of individuals, perhaps one could treat the domain of individuals in the model as the quantifier domain restriction supplied by context. We call this approach, the *model theoretic* approach.

The model theoretic approach works very well for sentences which contain only one quantified expression, such as:

(30) Everyone smokes.

According to the model theoretic approach, the (informal) truth-clause for (30) is:

(31) “Everyone smokes” is true relative to a model  $M$  iff everyone in the domain of  $M$  smokes in  $M$ .

Suppose that the domain of quantification for a particular utterance of (30) is the students in Mr Desiato’s third-grade class. We incorporate this fact into the semantic theory, according to the Model Theoretic approach, by considering the truth of (30) relative to models in which the domain is the set of students in Mr Desiato’s

<sup>37</sup> One might think that the variable-free approach so clearly presented in Cresswell (1996) is a metalinguistic approach. Not so. On this approach, quantifier domain restriction is treated via the provision to Logical Forms of a series of covert, object-language *operators*.

third-grade class. The sentence is true in such models just in case every member of the domain smokes, and false otherwise.

Unfortunately, the model theoretic approach fails for sentences containing more than one quantified expression. Consider, for example:

(32) Every sailor waved to every sailor.

This sentence can express the proposition that every sailor on the ship waved to every sailor on the shore (cf. Stanley and Williamson (1995)). The moral of such examples is, as Scott Soames (1986, 357) has written, “that contextual supplementation works at the level of constituents of sentences or utterances, rather than the level of the sentences or utterances themselves” (cf. also Westerståhl (1985), and Recanati (1996)). The model theoretic approach, in contrast, associates whole sentences with domains of quantification. This is simply not fine-grained enough to capture the phenomenon of quantifier domain restriction.<sup>38</sup>

However, the model theoretic approach is not the only metalinguistic semantic parameter approach. One can associate domains with quantifier expressions metalinguistically, in such a way as to incorporate Soames’s moral. To do so, one must provide some way for the context to shift within a clause.<sup>39</sup> Given that context can shift within a clause, one can treat examples such as (32) by providing the following sort of meaning rule for lexical items such as “sailor”:

(33) The denotation of “sailor” in *c* is the set of sailors in the domain provide by *c*.

If context can shift within a clause, then different sets can be provided as domains for the two occurrences of “every sailor” in (32). The set provided by the context relative to which the first occurrence is evaluated can be the set of things on the ship, and the

<sup>38</sup> This is not to say that the model theoretic approach is of no use in addressing problems of context dependence. One way to use models to reflect the effects of context is in dealing with the sort of examples raised in Reimer (1998). According to Reimer, there are cases in which context forces us to ignore certain *facts*. In Reimer’s cases, to capture the intuitive truth of certain utterances, one has to ignore the fact that the individuals under consideration have certain properties. One can account for these examples by treating the effect of context as restricting the evaluation of the utterances under consideration to models in which the individuals in question lack those properties; in effect treating the models as possible situations in which some actual facts do not obtain.

However, Reimer’s examples are still problematic for some standard approaches. For example, Davidsonian approaches operate with a non-relativized notion of truth; that is, a definition of truth that is not relative to a model. On a Davidsonian approach, it would therefore appear as if one would have to deny the semantic significance of the intuitions at issue in the examples. However, it is unclear that this poses a problem for Davidsonian approaches. Appeal to an absolute notion of truth might provide just the sort of principled basis for rejecting the significance of the intuitions that Reimer claims is difficult or impossible to produce (cf. Reimer (1998, 100–3)).

<sup>39</sup> For a treatment of demonstratives and indexicals along these lines, see the appendix to Braun (1996).

set provided by the context relative to which the second occurrence is evaluated can be the set of things on the shore.

Though this version evades some of the difficulties facing the model-theoretic approach, it too is ultimately unsatisfactory. Furthermore, it is unsatisfactory for reasons that apply to any version of a metalinguistic approach. In the last section, we discussed the phenomenon of quantified contexts, in which a domain of a quantified expression varies with the values introduced by a higher operator. This phenomenon is amenable to treatment in an approach in which logical forms contain variables that are assigned quantifier domains, for on such approaches, the quantifier domain variable is then bound by the higher operator. However, the phenomenon of quantified contexts is not amenable to treatment in a metalinguistic approach to quantifier domain restriction, since there is no variable to be bound.

Simplifying dramatically for the purposes of exposition (including omitting reference to “exactly”), we can consider the syntax of (24) to be as in:

- (34) [S<sub>PP</sub>[P In] [NP<sub>i</sub> [DET most] [N' [PP [P of] [NP John's Classes]]]]]  
 [S [NP<sub>i</sub> he] [VP [V fails] [NP [DET three] [N' [N Frenchmen]]]]].

The natural interpretation of (34) is that, where *x* is a class of John's, John fails three Frenchmen in *x*. To capture this reading, we need to postulate a variable bound by “most of John's classes” which is in some way associated with the quantifier phrase “three Frenchmen”. The different values of these variables correspond to the different quantifier domains for the phrase “three Frenchmen”. So, the phenomenon of quantified contexts shows that quantifier expressions contain variables which can be bound, and whose purpose is to supply the domains for quantifier expressions.

However, it is not completely straightforward how, in examples such as (24)–(26), quantifier domains vary as a function of the values introduced by the preceding quantifiers. Consider example (24). Here, the quantifier domain for “three Frenchmen” varies as a function of the values introduced by “most of John's classes”. But the values introduced by this latter quantifier expression are classes (in the educational sense), which are not appropriate entities to be quantifier domains. If quantifier domains are sets, then the context must provide a function from classes (in the educational sense) to sets. This function will map a class onto the set of students in that class. Similarly, to treat example (26), context must provide a function from events to sets. A successful analysis of these constructions must incorporate such a function (see von Stechow (1994, section 2.2.2), and Cresswell (1996, 81–7)).

What the phenomenon of quantified contexts shows is that, assuming standard treatments of binding, quantifier domain restriction must be treated with the use of variables. But we have to be somewhat subtle about the use of variables here. We cannot simply replace the noun “bottle” in (1) with a variable in logical form, since common nouns are obviously not pure demonstrative or indexical expressions. We must rather

associate variables with some part of quantifier expressions, such as “every bottle”. Distinct non-metalinguistic treatments of the problem of quantifier domain restriction correspond to different ways of associating variables with quantifier expressions.

However, we first need to say more about what it is to *associate* a variable with an expression. In our terminology, a variable is associated with a syntactic element, either a syntactic category or a lexical item, just in case it *co-habits* a node with it. We represent the co-habitation relation as an ordered pair of an expression and a variable. The details of this treatment will emerge in our discussion of the different accounts of quantifier domain restriction.

We first present our favored way to use contextual variables to treat the problem of quantifier domain restriction. Suppose each common noun (e.g., “bottle” and “cat”) co-habits a node with a contextual variable.<sup>40</sup> Then, a sentence such as:

(35) Every man runs

would receive the representation:




Figure 2

The value of “*i*” is an object provided by the context, and the value of “*f*” is a function provided by the context that maps objects onto quantifier domains. The restriction on the quantified expression “every man” in (35), relative to a context, would then be provided by the result of applying the function that context supplies to “*f*” to the object that context supplies to “*i*”.

By representing quantifier domain variables as co-habiting nodes with common nouns, we do not mean to deny to them syntactic reality. If binding is ultimately a syntactic relationship between a binder and a variable, then our arguments in the previous section demonstrate that they do indeed exist in the syntax (see pp. 46–64). Indeed, in what follows, we use the assumption that quantifier domain variables are syntactically real in some of our arguments against alternative accounts.

<sup>40</sup> The use of contextual variables as parasitic on other syntactic elements is an idea we have adopted from Westerståhl’s classic (1985). Westerståhl’s approach, however, differs in important respects from ours. According to him, contextual variables co-habit nodes with the categorial label, “Det”. This analysis is a mixture of the second and third alternatives to our approach, discussed below, and inherits the problems of both.



Before we give an explicit semantics, we note a simplifying assumption. In what follows, we treat quantifier domains as *sets*. However, the domains contexts provide for quantifiers are better treated as intensional entities such as *properties*, represented as functions from worlds and times to sets. Here is why. If quantifier domain restriction is a semantic process of a variable being assigned a value by the context, then, following standard semantic procedure, these values will be assigned before the sentence is evaluated with respect to other possible worlds or times. It follows that, relative to a context, contextual variables *rigidly designate* their values. However, relative to other worlds and times, the set corresponding to the quantifier domain may vary.

Suppose that John has a strange habit of buying exactly 70 bottles every time he goes to a supermarket. Suppose that John visits a supermarket that has exactly 70 bottles on the shelf, and purchases every bottle. Someone could then truly utter the sentence:

- (37) If there were a few more bottles on the shelf, John would not have purchased every bottle.

However, if we assign to the contextual variable associated with “every bottle” the set of bottles in the supermarket in the context of utterance of John’s sentence, given the standard semantics for counterfactuals, (37) could not be truly uttered. To capture the reading of (37) on which it is true, one must treat the entity assigned to the contextual variable as a function from worlds and times to, say, the sets of bottles in the relevant supermarket at those worlds and times.

However, treating quantifier domains as properties only would complicate the discussion. All of the points we make in this section remain valid even when the extra complexities attached to the addition of intensional operators are factored in.

Let us now turn to the semantics. We adopt the by now standard generalized quantifier treatment of quantifiers such as “every”, whereby they express relations between sets (cf. Barwise and Cooper (1981); Westerståhl (1989)). On this account, the semantic clauses for quantifiers such as “every” and “some” are as in (38):

- (38) (a) Every A B iff  $A \subseteq B$ .  
 (b) Some A B iff  $A \cap B \neq \emptyset$

On this account, the first argument of a quantified expression is determined by the head noun, and the second argument is determined by the verb phrase. For example, in the case of a sentence such as (35), the first argument would then be the set of men, and the second argument would be the set of runners.

No adjustment is required to extend the standard generalized quantifier treatment to interpret structures such as (36). But we do need to say something about the interpretation of expressions such as “⟨man, f(i)⟩”. Since we are taking quantifier domains to be sets, relative to a context, what results from applying the value of “f”

to the value of “i” is a set. Relative to a context, “f” is assigned a function from objects to sets. Relative to a context, “i” is assigned an object. The denotation of “⟨man, f(i)⟩” relative to a context *c* is then the result of intersecting the set of men with the set that results from applying the value given to “f” by the context *c* to the value given to “i” by *c*. That is (suppressing reference to a model to simplify exposition), where “[ $\alpha$ ]<sub>*c*</sub>” denotes the denotation of  $\alpha$  with respect to the context *c*, and “*c*( $\alpha$ )” denotes what the context *c* assigns to the expression  $\alpha$ :

$$(39) [\langle \text{man}, f(i) \rangle]_c = [\text{man}] \cap \{x: x \in c(f)(c(i))\}$$

In the case of (36), the resulting set is then the first argument of the generalized quantifier “every”.<sup>41</sup>

It is important, in understanding our analysis, to remember that what co-habits a node with common nouns are variables of the form “f(x)”, where “f” is a variable that is assigned a function from objects to sets by contexts, and “x” is a variable that is assigned objects relative to contexts. However, to simplify exposition in what follows, we will represent quantifier domain variables more simply as just single variables.

There is not much to disagree with in the claim that our favored semantic treatment is *one* possible way to accommodate a domain variable into the logical form of a sentence like (35). But why choose this semantic theory over others? In what follows, we will try to justify our preference. To do so, we will discuss three alternatives to our proposal. According to the first, domain variables co-habit *non-terminal nodes* with other syntactic elements. For example, the domain variable in the logical form for (35) co-habits a non-terminal node, namely the one labeled with *NP* dominating both the terminal node labeled with *every* and the terminal node labeled with *man*:




Figure 3

<sup>41</sup> One might think that one has to relativize both arguments of the quantifier to a domain. Not so. It is generally accepted that all true quantifiers found in natural language have the following two properties. First, they are *conservative*; that is, where *Q* is a binary natural language quantifier,  $QAB \leftrightarrow QA (A \cap B)$ . Secondly, they satisfy the *extension principle*; that is, where  $A, B \subseteq D \subseteq D'$ , then, where *Q* is a binary quantifier, *QAB* is true relative to the domain of quantification *D* if and only if *QAB* is true relative to the domain of quantification *D'*. This entails that the result of relativizing the first argument of a natural language quantifier to a domain of quantification *D* is equivalent to the result of relativizing both arguments to *D* (for the proof, see Westerståhl (1985, 54–5); for discussion, see Keenan and Westerståhl (1997, 852 ff.)).

According to the second, the quantifier domain variable occupies its own terminal node in the logical forms of sentences containing quantifier expressions. For example, the logical form of (35), on this approach, would be:

(41)




Figure 4

Finally, according to the third approach, the domain variable co-habits a terminal node with the quantifier “every”, rather than one with the noun “man”. On this account, the logical form of (33) would be as in (42):

(42)




Figure 5

According to the first alternative, contextual variables occur in non-terminal nodes. To interpret the resulting structures, such as (40), one needs the following sorts of composition rules (we give these composition rules in basic Montague Grammar). Let “[ $\alpha$ ]<sub>M,c</sub>” denote the denotation of  $\alpha$  with respect to the model M and context c. As is standard, verb phrase denotations and noun denotations are functions from objects to truth-values; noun-phrase denotations are functions from such functions to truth-values; and determiner denotations are functions from functions from objects to truth-values to noun-phrase denotations.

- (43) (a) [<sub>s NP<sub>i</sub> VP</sub>]<sub>M,c</sub> = t iff [<sub>NP<sub>i</sub></sub>]<sub>M,c</sub> ([<sub>VP</sub>]<sub>M,c</sub>) = t.
- (b) [<sub>NP<sub>i</sub> Det N</sub>]<sub>M,c</sub> = [<sub>Det</sub>]<sub>M,c</sub> ([<sub>N</sub>]<sub>M,c</sub> ∩ c(i)).

It is the composition rule (43b) which we find objectionable. What it does is apply the function denoted by the determiner to the result of intersecting the denotation of the noun with the value of the contextual variable. Our worry with composition rules of this sort is that they violate *compositionality*. Compositionality can be stated as follows: Suppose  $\alpha$  is a non-terminal node immediately dominating  $\beta_1 \dots \beta_n$ . Then there is a function  $f$  such that  $[\alpha] = f([\beta_1], \dots, [\beta_n])$ .<sup>42</sup>

According to (43b), the denotation of a noun phrase consisting of a determiner and a noun is determined by more than just applying the function expressed by the determiner to the denotation of the noun. Rather, in different contexts, the denotation of the noun phrase is a different function of the denotation of the determiner and the denotation of the head noun. Therefore, (43b) violates compositionality. One sometimes hears the view, voiced by philosophers of a pessimistic persuasion, that context dependence poses a serious worry for the project of giving a systematic natural language semantics. This view is sometimes expressed by the thesis that context dependence threatens the possibility of a compositional semantics. If composition rules of the kind in (43b) were required to accommodate the influence of context on semantic interpretation, then, on the assumption that compositionality is required for systematic semantics, there would be reason to adopt this pessimistic viewpoint.

We are not ourselves completely convinced that systematic semantics requires the truth of compositionality. For example, several semanticists have argued that full-blown compositionality should be rejected, though the sorts of potential failures they discuss would not undermine the systematicity of semantics. However, the sort of failures contemplated by these authors involve the alleged dependence of semantic value on *linguistic* context.<sup>43</sup> The sort of failure of compositionality involved in rules of the kind in (43b) is significantly more drastic than this. Rules such as (43b) allow the composition rules associated with a single linguistic structure to vary freely as a function of *extra-linguistic* context. An indefinite number of composition rules that vary not just according to linguistic context, but also according to extra-linguistic context, seems in tension with learnability considerations. If contextual variables were allowed to occur in non-terminal nodes, then the semantic theories required to interpret the resulting structures would violate compositionality in this quite drastic

<sup>42</sup> For expository purposes, we suppress reference to a model and a context. It is worth mentioning that this version of compositionality is rather lax, since it allows different syntactic configurations to involve different modes of semantic composition. According to some versions of compositionality (e.g. Lewis (1970); Heim and Kratzer (1998)) compositionality is the stronger requirement that there is only one mode of semantic composition.

<sup>43</sup> One example of this line of thought is found in discussions of game-theoretical semantics, which is not compositional (e.g. Jaako Hintikka and Gabriel Sandu (1997, 370 ff.)). Another author who rejects compositionality for reasons involving the dependence of semantic value on linguistic context is James Higginbotham (see his discussion of the "Indifference Principle" (1986, 33 ff.)).

manner. Therefore, the distribution of contextual variables should be restricted to terminal nodes. Our favored treatment of quantifier domain restriction accords with this restriction on the distribution of contextual variables, whereas the postulation of structures such as (40) does not.

But what about the second and third alternatives to our favored approach? These do not drastically violate compositionality, so we need other arguments to show that they are theoretically inferior to our proposal. Consider the second proposal, according to which contextual variables occupy their own terminal nodes. We do not have a decisive objection against this account. But it does seem to us to be methodologically troublesome in the following respect. According to it, quantifier domain variables occupy their own terminal nodes. This involves the postulation of an entire unarticulated relative clause. Such a postulation ultimately requires syntactic justification.

Our worry is not that such a syntactic justification is impossible to provide. It is rather that, without compelling reasons, one should not place such a burden on syntactic theory. We do think that occasionally semanticists have such reasons. But we do not see that this is so in the case of this approach to quantifier domain restriction. For there is a perfectly legitimate theoretical alternative which does not postulate extra syntactic complexity, namely the one we advocate. It is therefore to be preferred over one that does.

Our concern with the third alternative arises from a consideration of facts involving cross-sentential anaphora. Consider the sentence:

(44) Most people regularly scream. They are crazy.

Suppose the domain is the set of things in a certain village. There are two anaphoric readings of the pronoun in the second sentence in (44). On the first reading, it refers to all of the people in the village. On the second reading, it refers to those people in the village who regularly scream. The third alternative has difficulties explaining both of these readings.

Consider the first reading, that everyone in the village is crazy. Ideally, one would wish to say that cross-sentential anaphora of this sort requires antecedents that are constituents (nodes) of a preceding logical form. However, if the domain variable co-habits a terminal node with “most”, there is no single node in the logical form of the first sentence of (44) whose associated semantic value is the set of people in the village. In our favored approach, however, there is such a node: the one labeled with “⟨people, i⟩”. So, our favored approach provides a far more natural account of the first reading of the second sentence of (44).

Our favored approach also does a much better job with the second reading of the second sentence of (44). One theory in the literature which elegantly captures this reading is presented by Stephen Neale. He treats “they” as a proxy for a

certain description reconstructable from the logical form of the first sentence. The reconstruction is guided by the following principle:

If  $x$  is a pronoun that is anaphoric on, but not  $c$ -commanded by a non-maximal quantifier “[ $Dx:Fx$ ]” that occurs in an antecedent clause “[ $Dx:Fx$ ]( $Gx$ )”, then  $x$  is interpreted as “[the  $x$ :  $Fx \& Gx$ ]”.<sup>44</sup>

But note what happens if we apply this rule to (44), on the assumption that the domain variable co-habits the node of the quantifier “most”. In constructing the definite description which gives the interpretation of the pronoun in the second sentence, we drop the quantifier and lose the domain variable with it. “They” is then interpreted as [the  $x$ : person ( $x$ ) & regularly-scream( $x$ )], which results in an unrestricted reading of the second sentence, according to which everyone in the universe who regularly screams is crazy. By contrast, if, as on our favored approach, the variable co-habits a node with “person”, “they” is interpreted as [the  $x$ : (person,  $i$ ) ( $x$ ) & regularly-scream( $x$ )], we obtain the desired reading of the second sentence, according to which it expresses the proposition that every person in the village who regularly screams is crazy.<sup>45</sup>

In previous sections, we located the source of the solution to the descriptive problem of context dependence for quantifier expressions in the semantics. In this section, we have advanced and defended a particular semantic approach. We thereby take ourselves to have provided a satisfactory solution to the descriptive problem of context dependence for quantifier expressions. A sentence such as (1) can *communicate* a proposition concerning a restricted domain of bottles, because, relative to certain contexts, it *expresses* such a proposition. It expresses such a proposition relative to certain contexts because common nouns such as “bottle” always occur with a domain index. It follows that, in the logical form of quantified sentences, there are variables whose values, relative to a context, are (often restricted) quantifier domains.

<sup>44</sup> Neale (1990, 266, rule (P5b)). A quantifier “[ $Dx:Fx$ ]” is non-maximal iff there is some  $G$  for which “[ $Dx:Fx$ ]( $Gx$ )” is true but “[every  $x$ :  $Fx$ ]( $Gx$ )” is false. So, “most” is non-maximal.

<sup>45</sup> There are three main alternatives to Neale’s account (which is a development of that of Evans (1977)). First, there are DRT-type theories (for a DRT treatment of plural cross-sentential anaphora, see Kamp and Reyle (1993, 345–56)). Secondly, there is the “context-dependent quantifier” approach of Jeffrey King (1987, 1994). The point we made carries over to these two frameworks with some modifications. For example, on King’s theory, expressions that are anaphoric on non-symmetric quantifiers like “most” do not inherit the quantificational force of their antecedent. So, it is difficult to see how to transfer to the interpretation of the anaphoric expression domain variables that would occur on such quantifiers. Finally, there are pragmatic style solutions. These resemble Neale’s account except that they do not attempt the reconstruction of the antecedent clause from the LFs of the previous sentences. Such an account is advocated in Cooper (1979); Heim (1990); and Chierchia (1995). However, even such approaches should grant that the default interpretation of such pronouns is the semantic value of the expression given by a reconstruction algorithm of the sort found in Neale. The assumption that an algorithm like Neale’s provides a default semantic value is strong enough for our purposes.

## References

- ACZEL, P., ISRAEL, D., KATAGIRI, Y., and PETERS, S. (eds.) (1993), *Situation Theory and its Applications* (Stanford: Stanford University Press).
- ALMOG, J., PERRY, J., and WETTSTEIN, H. (eds.) (1989), *Themes from Kaplan* (Oxford University Press).
- BACH, K. (1992), 'Intentions and Demonstrations', *Analysis*, 52, 140–6.
- BACH, K. (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BACH, K. (1999), 'The Semantics Pragmatics Distinction: What It Is and Why It Matters', in Turner (1999), 65–84.
- BARWISE, J., and COOPER, R. (1981), 'Generalized Quantifiers and Natural Language', *Linguistics and Philosophy*, 4: 159–219.
- BARWISE, J., and COOPER, R. (1993), 'Extended Kamp Notation: A Graphical Notation for Situation Theory', in Aczel, Israel, Katagiri, and Peters (eds.) (1993), 29–53.
- BLACKBURN, W. K. (1998), 'Wettstein on Definite Descriptions', *Philosophical Studies*, 53: 263–78.
- BRAUN, D. (1994), 'Structured Characters and Complex Demonstratives', *Philosophical Studies*, 74: 193–219.
- BRAUN, D. (1996), 'Demonstratives and their Linguistic Meanings', *Nous*, 30: 145–73.
- CHIERCHIA, G. (1995), *Dynamics of Meaning* (Chicago: University of Chicago Press).
- CHOMSKY, N. (1995), *The Minimalist Program* (Cambridge, Mass.: MIT Press).
- COLE, P. (ed.) (1978), *Syntax and Semantics*, 9 (New York: Academic Press).
- COOPER, R. (1979), 'The Interpretation of Pronouns', in Heny and Schnelle (eds.) (1979), 61–92.
- COOPER, R. (1993), 'Generalized Quantifiers and Resource Situations', in Aczel, Israel, Katagiri, and Peters (eds.) (1993), 191–211.
- CRESSWELL, M. (1973), *Logics and Languages* (London: Methuen).
- CRESSWELL, M. (1996), *Semantic Indexicality* (Dordrecht: Kluwer).
- CRIMMINS, M. (1992), 'Context in the Attitudes', *Linguistics and Philosophy*, 15: 185–98.
- DALRYMPLE, M., SHIEBER, S., and PEREIRA, F. (1991), 'Ellipsis and Higher-order Unification', *Linguistics and Philosophy*, 14: 399–452.
- EVANS, G. (1977), 'Pronouns, Quantifiers, and Relative Clauses (I)', *Canadian Journal of Philosophy*, 7: 467–536.
- GABBAY, D., and GUENTHNER, F. (eds.) (1989), *Handbook of Philosophical Logic*, Vol. 4 (Dordrecht: Kluwer).
- GARCIA-CARPINTERO, M. (1998), 'Indexicals as Token-reflexives', *Mind*, 107: 529–63.
- GAUKER, C. (1997), 'Domain of Discourse', *Mind*, 106: 1–32.
- GAUKER, C. (1998), 'What is a Context of Utterance?', *Philosophical Studies*, 1–24.
- GROENENDIJK, J. (ed.) (1981), *Formal Methods in the Study of Language* (Amsterdam: Mathematisch Centrum).
- GROENENDIJK, J., and STOKHOF, M. (1991), 'Dynamic Predicate Logic', *Linguistics and Philosophy*, 14: 39–100.
- GUNDERSON, K. (ed.) (1975), *Minnesota Studies in the Philosophy of Science*, Vol. 7 (Minneapolis: University of Minnesota Press).
- HALE, B., and WRIGHT, C. (eds.) (1997), *A Companion to the Philosophy of Language* (Oxford: Blackwell).

- HEIM, I. (1982), *The Semantics of Definite and Indefinite Noun-phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- HEIM, I. (1990), 'E-type Pronouns and Donkey Anaphora', *Linguistics and Philosophy*, 13: 137–78.
- HEIM, I., and KRATZER, A. (1998), *Semantics in Generative Grammar* (Oxford: Blackwell).
- HENY, F., and SCHNELLE, H. (eds.) (1979), *Syntax and Semantics*, 10 (New York: Academic Press).
- HIGGINBOTHAM, J. (1986), 'Linguistic Theory and Davidson's Program in Semantics', in LePore (1986), 29–49.
- HINTIKKA, J., and SANDU, G. (1997), 'Game Theoretical Semantics', in van Benthem and ter Meulen (eds.) (1997), 361–410.
- HIRST, G. (1987), *Semantic Interpretation and the Resolution of Ambiguity* (Cambridge University Press).
- ISRAEL, D., and PERRY, J. (1996), 'Where Monsters Dwell', in Seligman and Westerståhl (eds.) (1996), 303–16.
- KAMP, H. (1981a), 'The Paradox of the Heap', in Mönnich (1981), 225–77.
- KAMP, H. (1981b), 'A Theory of Truth and Semantic Representation', in Groenendijk (1981).
- KAMP, H., and REYLE, U. (1993), *From Discourse to Logic* (Dordrecht: Kluwer).
- KAPLAN, D. (1978), 'Dthat', in Cole (1978), 221–43.
- KAPLAN, D. (1989a), 'Demonstratives', in Almog, Perry, and Wettstein (eds.) (1989), 481–564.
- KAPLAN, D. (1989b), 'Afterthoughts', in Almog, Perry, and Wettstein (eds.) (1989), 565–614.
- KEENAN, E., and WESTERSTÅHL, D. (1997), 'Generalized Quantifiers in Linguistics and Logic', in van Benthem and ter Meulen (1997), 837–93.
- KING, J. (1987), 'Pronouns, Descriptions, and the Semantics of Discourse', *Philosophical Studies*, 51: 341–63.
- KING, J. (1994), 'Anaphora and Operators', in *Philosophical Perspectives*, 8: 221–50.
- KLEIN, E. (1985), 'VP Ellipsis in DR Theory', in van Benthem and ter Meulen (1985), 161–87.
- LEPORE, E. (1986), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (Oxford: Basil Blackwell).
- LEWIS, D. (1970), 'General Semantics', *Synthese*, 22: 18–67.
- LEWIS, D. (1975), 'Languages and Language', in Gunderson (1975), 3–35.
- MAY, R. (1985), *Logical Form: Its Structure and Derivation* (Cambridge, Mass.: MIT Press).
- MCGINN, C. (1981), 'The Mechanism of Reference', *Synthese*, 49: 157–86.
- MÖNNICH, U. (ed.) (1981), *Aspects of Philosophical Logic* (Dordrecht: Reidel).
- MONTAGUE, R. (1974), 'Pragmatics', in Montague, *Formal Philosophy* (New Haven: Yale University Press), 95–118.
- MUNITZ, M. K., and UNGER, P. K. (eds.) (1974), *Semantics and Philosophy* (New York: New York University Press).
- NEALE, S. (1990), *Descriptions* (Cambridge, Mass.: MIT Press).
- RECANATI, F. (1993), *Direct Reference: From Language to Thought* (Oxford: Blackwell).
- RECANATI, F. (1996), 'Domains of Discourse', *Linguistics and Philosophy*, 19: 445–75.
- REIMER, M. (1991), 'Demonstratives, Demonstrations, and Demonstrata', *Philosophical Studies*, 63: 187–202.
- REIMER, M. (1992), 'Three Views of Demonstrative Reference', *Synthese*, 93: 373–402.
- REIMER, M. (1998), 'Quantification and Context', in *Linguistics and Philosophy*, 21: 96–115.



- REYLE, U. (1993), 'Dealing with Ambiguities by Underspecification: Construction, Representation, and Deduction', *Journal of Semantics*, 10: 123–79.
- SALMON, N. (1986), *Frege's Puzzle* (Cambridge, Mass.: MIT Press).
- SCHIFFER, S. (1992), 'Belief Ascription', *Journal of Philosophy*, 499–521.
- SCHIFFER, S. (1995), 'Descriptions, Indexicals, and Belief Reports: Some Dilemmas (but not the ones you expect)', *Mind*, 104: 107–31.
- SELIGMAN, J., and WESTERSTÅHL, D. (eds.) (1996), *Logic, Language, and Computation*, Vol. 1. (Stanford: CSLI).
- SMITH, Q. (1989), 'The Multiple Uses of Indexicals', *Synthese*, 78: 167–91.
- SOAMES, S. (1986), 'Incomplete Definite Descriptions', *Notre Dame Journal of Formal Logic*, 349–75.
- SPERBER, D., and WILSON, D. (1986), *Relevance: Communication and Cognition* (Cambridge, Mass.: Harvard University Press).
- STALNAKER, R. (1973), 'Presuppositions', *Journal of Philosophical Logic*, 2: 447–57.
- STALNAKER, R. (1974), 'Pragmatic Presuppositions', in Munitz and Unger (1974), 197–214.
- STALNAKER, R. (1997), 'Reference and Necessity', in Hale and Wright (1997), 534–54.
- STANLEY, J., and WILLIAMSON, T. (1995), 'Quantifiers and Context-dependence', *Analysis*, 55: 291–5.
- TRAVIS, C. (1997), 'Pragmatics', in Hale and Wright (1997), 87–107.
- TURNER, K. (1999), 'The Semantics–Pragmatics Interface from Different Points of View', (Oxford: Elsevier).
- VAN BENTHEM, J., and TER MEULEN, A. (1985), *Generalized Quantifiers in Natural Language* (Dordrecht: Foris).
- VAN BENTHEM, J., and TER MEULEN, A. (eds.) (1997), *Handbook of Logic and Language* (Amsterdam: Elsevier).
- VAN DEEMTER, K., and PETERS, S. (eds.) (1996), *Semantic Ambiguity and Underspecification* (Stanford: CSLI).
- VON FINTEL, K. (1994), *Restrictions on Quantifier Domains*. Ph.D. thesis, University of Massachusetts, Amherst.
- WESTERSTÅHL, D. (1985), 'Determiners and Context Sets', in van Benthem and ter Meulen (1985), 45–71.
- WESTERSTÅHL, D. (1989), 'Quantifiers in Formal and Natural Languages', in Gabbay and Guenther (1989) 1–131.
- WETTSTEIN, H. (1981), 'Demonstrative Reference and Definite Descriptions', *Philosophical Studies*, 40: 241–57.
- WETTSTEIN, H. (1984), 'How to Bridge the Gap Between Meaning and Reference', *Synthese*, 58: 63–84.

# 3

## Nominal Restriction

Extra-linguistic context appears to have a profound effect on the determination of what is expressed by the use of linguistic expressions. For a bewildering range of very different linguistic constructions, adhering to relatively straightforward linguistic intuition about what is expressed leads us to the conclusion that facts about the non-linguistic context play many different roles in determining what is said. Furthermore, that so many different constructions betray this sort of sensitivity to extra-linguistic context understandably leads to pessimism about rescuing the straightforward intuitions while preserving any sort of systematicity in the theory of meaning.

A presumption motivating the pessimistic inclination is that, if we accept the ordinary intuitions, what *appear* to be very different ways in which context affects semantic content in fact *are* different ways in which context affects linguistic content. Pessimism is a natural reaction to those who adopt this presumption, because if appearance is a good guide to the facts in this domain, then there are just too many ways in which context affects semantic content to preserve systematicity. One common and natural reaction to these facts is, therefore, to deny the semantic significance of the ordinary intuitions, thereby relegating the project of explaining the apparent effects of extra-linguistic context on semantic content to a domain of inquiry outside the theory of meaning proper. So doing removes the threat context poses to the systematicity of semantic explanation, but at the cost of reducing the interest of the semantic project.

In this chapter, I explore a different reaction to the situation. My purpose is to undermine the presumption that what appear to be very different effects of context on semantic content are very different effects. My challenge is of necessity rather limited, since it is too implausible to trace all effects of extra-linguistic context on semantic content to the very same source. Rather, I will take, as a case-study, three superficially very different effects of context on semantic content, and show that they are due to the very same mechanism, what I call *Nominal Restriction*. I thereby hope to

Thanks to Richard Boyd, Herman Cappelen, Richard Heck, Kathrin Koslicki, Ernie Lepore, Kirk Ludwig, Brett Sherman, Anna Szabolcsi, and Zsófia Zvolensky for valuable comments and discussion. I would especially like to thank Delia Graff, Jeff King, and Zoltán Gendler Szabó for many hours of extraordinarily fruitful discussion about the topics of this chapter.

provide convincing evidence of the promise of the project of reducing all apparent effects of context on semantic content to a small number of sources.


In the first section, I introduce an account of the phenomenon of quantifier domain restriction presented in Chapter 2, and provide two novel defenses of it. In the second section, I turn to a discussion of comparative adjectives. As I argue, the theory introduced in the next section, which I call the *Nominal Restriction Theory*, also provides an explanation for some mysterious facts about how context determines the comparison class for uses of comparative adjectives. In the third section, I turn to another apparently very different sort of effect of extra-linguistic context on semantic content, and show how it too is smoothly explicable on the Nominal Restriction Theory. I then draw some consequences from the discussion for some issues in the theory of reference.

## I Domain Restriction

The sentence “Every bottle is empty” can be used to communicate many different propositions. For example, if John is about to go shopping, and is wondering whether he should buy something to drink, Hannah can utter “Every bottle is empty” to communicate the proposition that every bottle in the house is empty. In this section, I describe and defend what I believe to be the best account of how sentences containing quantified noun phrases such as “every” and “some” can be used to communicate propositions about a restricted domain of entities. In the rest of the chapter, I draw out some consequences of the account for other constructions.

The account I will defend is that presented on pp. 95–107. The simplest version is that each nominal expression is associated with a domain variable. Relative to a context, the domain variable is assigned a set. The semantic relation between the extension of the nominal expression and the set is set-theoretic intersection. A sentence such as “Every bottle is empty” can communicate the proposition that every bottle in Hannah’s house is empty, because, relative to the relevant context, the domain variable associated with “bottle” is assigned the set of things in Hannah’s house. “Every bottle is empty” communicates the proposition that every bottle in Hannah’s house is empty, because, relative to this context, it semantically expresses this proposition.


This is the theory in its simplest form. Details need to be filled in, and modifications added. To explain some of them, I will have to introduce some modest syntax. Let us call the output of the syntactic process that is visible to semantic interpretation a *logical form*. A logical form is a lexically and structurally disambiguated ordered sequence of word types, where word types are individuated both by semantic and syntactic properties. Logical forms are phrase markers. An example of such a phrase marker, for the sentence “Hannah loves Sue”, is as follows:



The nodes in this diagram are the points labeled either with syntactic categories or lexical items. So, “Hannah” labels a node, as does “N”, “VP”, and “V” (we shall also talk of labels of nodes as *occupying* these nodes). The nodes are connected by branches, which are the lines in the diagram. We say that a node *X dominates* another node *Y* in a phrase marker if there is a path of branches leading downward in the tree from *X* to *Y*. We say that a node *X* immediately dominates another node *Y* in a phrase marker just in case *X* dominates *Y*, and there is no node *Z* between *X* and *Y*. Nodes that dominate other nodes are called *nonterminal nodes*. Nodes that dominate no other nodes are *terminal nodes*. The nodes labeled with lexical items such as “Hannah” or “loves” are always terminal nodes. These are the objects that we assume are interpreted by semantic theories.

In the theory I have sketched, nominal expressions are associated with domain variables. By “association”, I mean that nominal expressions, such as “bottle”, cohabit a terminal node with a domain variable. So, on this account, the logical form of a sentence such as “Every man runs” is:

(I)



On this view, domain variables are independently meaningful expressions that incorporate with nouns. Thus, domain variables are not, on this account, expressions that occupy their own terminal nodes.

One reason to think that this is the right syntactic treatment of domain variables is that it is difficult to find sentences containing pronouns that are anaphoric on domain variables.<sup>1</sup> Similarly, expressions that are incorporated with other expressions often do not license anaphoric relations. For example, as Irene Heim (1982, 24) has pointed out, (2) easily allows a reading where the pronoun “it” is anaphoric on “bicycle”, whereas (3) does not:

<sup>1</sup> Thanks to Herman Cappelen for emphasizing to me the difficulty of finding natural examples of anaphoric dependence on a domain variable.

- (2) John owns a bicycle. He rides it daily.  
 (3) John is a bicycle-owner. He rides it daily.

Thus, the fact that it is difficult to have anaphora on domain variables is to be predicted, given representations such as (1).<sup>2</sup>

In the theory I have sketched, quantifier domain restriction is due to the presence of domain variables in the actual syntactic structure of sentences containing quantified noun phrases. But syntactic structure cannot simply be postulated on semantic grounds. Rather, evidence of a syntactic sort must be available for the existence of domain variables. The main source of syntactic evidence comes from the fact that domain variables interact in binding relations with quantifiers.

Here is the evidence from bindability. Consider the sentence:

- (4) Everyone answered every question.

(4) can express the proposition that everyone  $x$  answered every question on  $x$ 's exam. What this indicates is that there is a variable accessible to binding somewhere in the quantified phrase "every student". There are many other examples of this phenomenon, such as:

- (5) a. In most of his classes, John fails exactly three Frenchmen.  
 b. In every room in John's house, he keeps every bottle in the corner.

In all of these cases, the domain associated with a quantified noun phrase varies as a function of the values introduced by a previous quantified noun phrase. For example, (5a–b) intuitively mean something like:

- (6) a. In most of his classes  $x$ , John fails exactly three Frenchmen in  $x$ .  
 b. In every room  $x$  in John's house, he keeps every bottle in  $x$  in the corner.

On the assumption that binding is fundamentally a syntactic phenomenon, such examples provide evidence for a variable somewhere in the syntactic structure of quantified noun phrases.<sup>3</sup>

There is therefore syntactic evidence for the existence of domain variables in sentences containing quantifier expressions. But to treat examples such as (4) and (5a–b), however, the simple theory presented above must be modified. The quantifier "everyone" in (4) and the quantified expressions in (5a–b) range over objects. But quantifier domains are sets, rather than objects. For example, the quantifier domains associated with the quantified noun phrase "three Frenchmen" in (5a) are, for each class John

<sup>2</sup> Thanks especially to Tom Werner for discussion here.

<sup>3</sup> Farkas (1997) provides a sophisticated discussion of examples such as (4) and (5). However, Farkas does not share our assumption that binding is fundamentally a syntactic phenomenon. Instead, she provides her own semantic account of scope and binding, which she uses to explain the bound readings of (4) and (5).

teaches, the set of students in that class. Similarly, the quantifier domains associated with the quantified noun phrase, “every bottle” in (5b) are, for each room  $x$  in John’s house, the things that are in  $x$ . To reflect the kind of dependence at issue, we must adjust the syntax and semantics of quantified sentences. Instead of representations such as:



we require representations such as:



The value of “ $i$ ” is an object provided by the context, and the value of “ $f$ ” is a function provided by the context that maps objects onto quantifier domains. The restriction on the quantified expression “every man” in (7), relative to a context, would then be provided by the result of applying the function that context supplies to “ $f$ ” to the object that context supplies to “ $i$ ”.

Adopting the by now standard generalized quantifier treatment of quantifiers such as “every”, whereby they express relations between sets (cf. Barwise and Cooper (1981); Westerståhl (1989)), the semantic clauses for quantifiers such as “every” and “some” are as in (8):

- (8) a. Every  $A B$  iff  $A \subseteq B$ .  
 b. Some  $A B$  iff  $A \cap B \neq \emptyset$ .

On this account, the initial noun phrase determines the first argument of a quantified expression, and the second argument is determined by the verb phrase. For example, in the case of a sentence such as (7), the first argument would then be the set of men, and the second argument would be the set of runners.

No adjustment is required to extend the standard generalized quantifier treatment to interpret structures such as (7). But we do need to say something about the interpretation of expressions such as “ $\langle \text{man}, f(i) \rangle$ ”. Since we are taking quantifier domains to be sets, relative to a context, what results from applying the value of “ $f$ ” to the value of “ $i$ ” is a set. Relative to a context, “ $f$ ” is assigned a function from objects to sets. Relative to a context, “ $i$ ” is assigned an object. The denotation of “ $\langle \text{man}, f(i) \rangle$ ”

relative to a context  $c$  is then the result of intersecting the set of men with the set that results from applying the value given to “ $f$ ” by the context  $c$  to the value given to “ $i$ ” by  $c$ . That is (suppressing reference to a model to simplify exposition), where “[ $\alpha$ ] $_c$ ” denotes the denotation of  $\alpha$  with respect to the context  $c$ , and  $c(\alpha)$  denotes what the context  $c$  assigns to the expression  $\alpha$ :

$$(9) [\langle \text{man}, f(i) \rangle]_c = [\text{man}] \cap \{x: x \in c(f)(c(i))\}$$

In the case of (7), the resulting set is then the first argument of the generalized quantifier “every”. Here is how the theory works with sentences such as the ones in (4) and (5). Consider first:

(4') Everyone answered every question.

Intuitively, the interpretation of this sentence in the envisaged scenario is “everyone  $x$  answered every question on  $x$ 's exam”. According to the theory just outlined, “every question” is of the form “every  $\langle$  question,  $f(i) \rangle$ ”. The variable “ $i$ ” is bound by the higher quantifier “everyone”. Context supplies “ $f$ ” with a function from persons to the set of problems on that person's exam, yielding the desired interpretation.

Let us consider one more example, for instance:

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

Here is how the theory just sketched treats (5a). The intuitive interpretation of this sentence is “In most of his classes  $x$ , John fails exactly three Frenchmen in  $x$ ”. According to the theory just outlined, “three Frenchmen” is of the form “three  $\langle$  Frenchmen,  $f(x) \rangle$ ”. The variable “ $x$ ” is bound by the higher quantifier “most of his classes”. Context supplies “ $f$ ” with a function that takes a class and yields the set of students in that class. This set is then intersected with the set of Frenchmen, to yield the first argument of the generalized quantifier “three”.

Here is how the theory works with a simpler example. Suppose I say:

(10) Every fireman goes to Jack's bar.

Presumably, I intend to be speaking about firemen associated with a particular town or location; what I am asserting is that every fireman associated with location  $l$  goes to Jack's bar. On the theory just sketched, “every fireman” is of the form “every  $\langle$  fireman,  $f(i) \rangle$ ”. My intentions (for example) determine a location as the value of the variable “ $i$ ”. Furthermore, they determine a function from locations to sets of things that generally occupy those locations, which is the value of “ $f$ ”.<sup>4</sup> This function,

<sup>4</sup> In some cases, it may be implausible to suppose that speaker or participant intentions determine both an object and a function from objects to properties. In such cases, we may suppose that context supplies a set as value to “ $i$ ” and the identity function as a default value to the function variable “ $f$ ” (see p. 53 for a similar suggestion involving tense). Thanks to Brett Sherman for discussion here.

applied to the location  $l$ , yields the set of things generally occupying  $l$ . This set is then intersected with the set of firemen, to yield the first argument of the generalized quantifier “every”.


According to this theory of quantifier domain restriction, it is due to the fact that each nominal co-occurs with variables whose values, relative to a context, together determine a domain. Thus, if it is right, “quantifier domain restriction” is a misleading label; better would be “nominal restriction”. Accordingly, I will call this theory the Nominal Restriction Theory (NRT).

In the rest of this section, I will provide arguments in support of the controversial aspects of the NRT. There are two controversial properties of NRT. The first property is that quantifier domain indices are associated with common nouns such as “fireman” and “person”, rather than, as one might more naturally expect, determiners such as “every” and “some”. The second property is that quantifier domain restriction does not merely involve the contextual provision of a property or a set as the value of an element in the syntactic structure of quantified sentences. Rather, it involves the provision both of an object and a function from objects to sets to an individual variable and a function variable that occur along with every nominal expression. In the rest of the section, I will defend each of these two controversial commitments.

According to NRT, the intuitive restriction on quantificational determiners such as “every”, “some”, and “most” is not due, as may seem obvious, to a restriction on the quantificational expressions themselves, but rather to a restriction on the nominal complements of these determiners. This is an unintuitive feature of the theory, one that needs a justification. In Chapter 2, several arguments are advanced in support of the conclusion that quantifier domain variables occur with nominals rather than determiners. I will not repeat those arguments here. Rather, I want to present a different argument for the conclusion, due to Delia Graff (p.c.).

Suppose that the domain-restricting index occurred on the determiner, rather than the head noun. Here is how the syntax and semantics would work. Abstracting from the complexity involving the function variable, in this case, the structure of “Every man runs” would be:

(II)



instead of (1) or (7). The semantic clause for “every” would then be something like:

(12)  $\langle \text{Every}, i \rangle A B$  iff  $A \cap c(i) \subseteq B$ .



So, the semantic clause would intersect the first argument of the generalized quantifier with the set provided to *i* by context.

This theory works well for examples such as “Every man runs”. But Graff’s argument demonstrates that theory has problems with slightly more complex constructions, such as:

(13) The tallest person is nice.

Consider an occurrence of (13), relative to a context in which the domain comprises the students at Cornell University. Suppose further that the tallest student at Cornell is indeed nice. The proposition expressed by such an occurrence of (13) is true. But, on the account of domain restriction we are now considering, it is predicted to be false. Here is why.

Since no domain index is on “person”, its denotation will be the set of all people in the universe. We may suppose the semantic function of “tallest” is to select the tallest member of the denotation of the head noun. Suppose that Jan, a basketball player in Holland, is the tallest person in the world. So, the set that results from applying “tallest” to “person” is the singleton set containing Jan. Now, on this account of domain restriction, the semantic clause governing “the’ is:

(14)  $\langle \text{The}, i \rangle A B \text{ iff } |A \cap c(i)| = 1 \& (A \cap c(i)) \cap B = 1$

If we apply this theory to the relevant occurrence of (12), it follows that the first argument of the generalized quantifier “The” is the singleton set containing Jan, and the value of “*i*”, relative to the envisaged context, is the set of students at Cornell. But since Jan is not at Cornell, the theory of quantifier domain restriction in question predicts that, relative to the envisaged context, (13) is false (or perhaps truth-valueless).

It is clear what has gone wrong. In evaluating “tallest person”, one does not select the tallest person in the world. Rather, one selects the tallest person in the contextually relevant domain. But this demonstrates that the domain is associated with the head noun, rather than with the determiner.<sup>5,6</sup>

<sup>5</sup> One way to attempt to evade Graff’s argument is to formulate semantic rules that take the domain index on the determiner and intersect it with the nominal head of the complement of the determiner. Such rules are relatively straightforward to formulate. However, they violate in a quite drastic manner what Larson and Segal (1995, 78) call *strong compositionality*. Strong compositionality is the thesis that the interpretation of each node in a syntactic tree is a function only of the interpretation of the nodes it immediately dominates. Despite the name given to it by Larson and Segal, it is a principle that in fact is significantly weaker than the principle of compositionality that is presupposed in other textbooks (e.g., Heim and Kratzer (1998)).

<sup>6</sup> Graff’s argument presupposes that the superlative adjective is to be interpreted in its entirety *in situ*. Some linguists, however, give a treatment of superlative constructions in which the superlative operator “est” detaches from the adjective and either takes scope over the whole noun phrase, or incorporates with the determiner. Adoption of such a framework may perhaps undermine Graff’s argument.

The second controversial feature of NRT is that, according to it, quantifier domain restriction is due not to the provision of a set or property, but rather to the provision of an object and a function that yields a set or property for that object as value. Indeed, the account involves the postulation of a function variable in the logical form of quantified sentences, whose values, relative to contexts, are functions whose values are quantifier domains. As we have seen, there is evidence from bindability for the existence of object variables in the syntactic structure of quantified noun phrases. But one might wonder whether there is evidence of a non-semantic sort for the existence of the function variables postulated by NRT.<sup>7</sup>

How does one argue for the existence of a variable? One syntactic feature of variables is their capacity to be bound by quantificational expressions. So, one way to argue for the existence of a variable in a certain construction is to produce examples of that construction in which a higher operator binds the postulated variable. I have exploited this strategy already in arguing for the existence of object variables in quantified noun phrases. In what follows, I use it again in providing evidence for the existence of function variables in quantified noun phrases.

It is difficult to find natural examples in which the function variable that NRT postulates in nominal expressions is bound. The reason for the difficulty, of course, is that it is difficult to find natural language constructions that involve quantification over functions. One kind of construction that arguably does involve quantification over functions involves the so-called “functional reading” of questions. Here is an example of a functional reading of a question:

(15) Q: What does every author like?

A: Her first book.

On a “functional question” account of (15), its semantics involves quantification over functions. In particular, the question is interpreted as:

(16) What function  $f$  is such that for every author  $x$ ,  $x$  likes  $f(x)$ .

Interpreting the question in (15) in this manner is advantageous. For it allows us to capture the intuition that the answer to the question is a proposition concerning a function, while maintaining the standard semantics for questions, according to which the semantic value of a question is the set of its (contextually relevant) good answers.<sup>8</sup>

One might be skeptical of functional readings of questions, on the grounds that they are really just “pair-list” readings. (15), for example, also has a pair-list reading.

<sup>7</sup> I am grateful to Jeff King, in particular, for stressing the need for a non-semantic justification for the function variables postulated by NRT.

<sup>8</sup> One excellent discussion of functional readings is in Groenendijk and Stokhof (1984, ch. 3). Two other sources are Engdahl (1986) and Chierchia (1993).

On this reading, the answer to the question in (15) would be a list of authors and what they like, as in (17):

(17) Q: What does every author like?

A: Hannah likes her first book, Paula likes her first book, and Matt likes his first book.

One might think that there is no genuine semantic difference between functional readings and pair-list readings, and on this ground deny that (15) involves quantification over functions. However, there is good evidence against running together functional readings of questions with pair-list readings (cf. Groenendijk and Stokhof (1984, ch. 3)). For example, (18) has a functional reading, but no pair-list reading:

(18) Q: What does no author like?

A1: Her first book.

#A2: Hannah doesn't like her first book, Paula doesn't like her first book, and Matt doesn't like his first book.

In general, when a *wh*-phrase (a word such as “who”, “what”, or “which”) takes scope over a downward-monotonic quantifier, then pair-list readings are disallowed, but functional readings are not.<sup>9</sup> So, functional readings of questions simply cannot be assimilated to pair-list readings.

Here are two other arguments against assimilating functional readings to pair-list readings (from Elisabet Engdahl (1986, 167–8)). First, it would seem that (19) could be true, even if John is not acquainted with every author in the world:

(19) John knows which book every author in the world likes the least, namely her first.

However, if we suppose that functional readings were a special case of pair-list readings, then (19) could not be true in this situation. For then (19) would be equivalent to:

(20) John knows, of every author in the world, which book she likes the least.

But (20) is clearly false if John is not acquainted with every author in the world.

Secondly, if functional readings were a special case of pair-list readings, then one would expect (21) to be contradictory:

(21) John knows which woman every Englishman admires most, namely his mother, but he doesn't know who the women in question are.

But (21) (Engdahl's (39)) does not seem contradictory; indeed, it could very well be true. But if functional readings were a special case of pair-list readings, then (21) should be contradictory.

<sup>9</sup> A quantifier *Q* is downward monotonic if and only if, where *QA*, then for all *B* such that  $B \subseteq A$ , *QB*.

These three arguments together provide powerful evidence in favor of the thesis that questions admit of a distinct functional interpretation. But, as I now demonstrate, if there are such readings, then there are clear examples in which the function variables postulated in nominal expressions by NRT are bound.

Suppose that John and Bill are arguing about which branch of the armed forces is the best. Bill has been arguing that the Navy is superior. John, an advocate of the marines, somewhat rhetorically asks:

(22) But what is every person truly proud to belong to?

where the intuitive answer is supposed to be “the marines”. So, (22) has the interpretation:

(23) What  $x$  is such that every person  $y \in x$  is truly proud to belong to  $x$ ?

It is a functional reading of a certain variation of examples such as (22) that will provide us with the evidence for a bound reading of the function variable postulated by the account of domain restriction outlined above. One example of this kind is (24):

(24) Q: In every country, what is every person proud to belong to?  
A: Its Progressive Party.

The interpretation of the question in (24) is:

(25) In every country  $c$ , what function  $f$  is such that every person  $x \in f(c)$  is proud to belong to  $f(c)$ ?

(24), so interpreted, is one example that provides evidence for the syntactic reality of the function variable postulated in nominal expressions by NRT. For according to NRT, “person” co-occurs with a function variable “ $f$ ”, and an object variable “ $x$ ”. The interpretation given in (25) is one in which the object variable associated with “person” is bound by the quantifier “every country”, and the function variable is bound by “what”, which here has the force of an existential quantifier over functions.

Of course, this example is rather complex. But the reason it is complex is simply because it is difficult to find examples of natural language constructions involving quantification over functions. It is striking that the linguistically most compelling case of such quantification provides straightforward examples in support of NRT.

The two most controversial properties of NRT are, first, that quantifier domain indices are associated with nominal expressions rather than with quantificational determiners, and second, that it postulates function variables in the syntactic structure of sentences containing quantified noun phrases. In this section, I have argued that both of these properties are independently motivated. If NRT is correct, then quantifier domain restriction is an effect of nominal restriction. But quantifier

domain restriction is not the only effect of nominal restriction. In the next two sections, I want to explore other effects of nominal restriction. This will complete my argument for the thesis that many superficially distinct kinds of dependence of semantic value on context are due to the same source.

## II Adjectives and Comparison Classes

There are several ways in which sentences containing comparative adjectives, such as “small”, “tall”, “heavy”, and “large” are sensitive to context. One salient way involves the provision of a comparison class. Consider predicative uses of a comparative adjective, such as:

- (26) That building is small.
- (27) That basketball-player is short.
- (28) That flea is small.

On one natural reading of (26), the building in question is not being said to be small for an object in general (whatever that may mean). Rather, the building is being said to be small for a building. Similarly, on a natural reading of (27), the basketball-player in question is not being said to be small for a person, but only for a basketball-player. Finally, (28) shows that there is an equally natural reading of these constructions in which the comparison class is not provided by the sentence. For, on a natural reading of (28), what it expresses is that the flea in question is small for an animal.

On what is perhaps the classical account of predicative uses of comparative adjectives (e.g. Parsons (1972, 139); Siegel (1975, 26–8)), the sentences (26)–(28) on these interpretations, are elliptical for (29–31):<sup>10</sup>

- (29) That building is a small building.
- (30) That basketball-player is a short basketball-player.
- (31) That flea is a small animal.

On this account, the context-sensitivity of (26)–(28) is resolved by postulating logical forms in which the comparison class is the denotation of a nominal expression provided by context.

<sup>10</sup> (29)–(31) are themselves ambiguous between what is sometimes called the intersective reading and the non-intersective reading (for a useful discussion, cf. section 1.0 of Larson (1998)). For example, “Mugsy is a short basketball-player” can mean either that Mugsy is short for a basketball-player (the non-intersective reading) or (for example) that Mugsy is short for a human being and Mugsy is a basketball-player. According to the view I am defending, (26)–(28) are elliptical for (29)–(31), in their *non-intersective* use. I will not address the intersective readings of sentences such as (29)–(31).

However, many linguists reject this classical account of predicative uses of comparative adjectives. Perhaps the central reason for abandoning the classical account is that, as Hans Kamp has written, in arriving at a comparison class, “the noun is not always the only determining factor” (Kamp (1975, 152)). Consider, for example:

- (32) Smith is a remarkable violinist.  
 (33) Fred built a large snowman.

As Kamp (1975, 152–3) notes, (32) may be true “when said in comment on his after-dinner performance with the hostess at the piano, and false when exclaimed at the end of Smith’s recital in the Festival Hall—even if on the second occasion Smith played a bit better than on the first”. Similarly, suppose that Fred is a 7-year-old child. An occurrence of (33) may still be true, if Fred has built a snowman that is large for a snowman built by a 7-year-old child. That is, the comparison class for “large” in (33) is not just given by the denotation of the nominal “snowman”. Similarly, the comparison class for “remarkable” in (32) is not just given by the nominal “violinist”. Rather, the comparison classes are considerably narrower than the extensions of these nominals.

According to the classical account of predicative uses of comparative adjectives, they are elliptical for constructions in which the nominal complement of the comparative adjective is present. However, providing a nominal does not yet specify the comparison class. Furthermore, one might expect that any satisfactory account of this “extra” context-dependency would be up to the task of supplying the entire comparison class on its own, without postulating a hidden nominal. Therefore, it seems, this traditional account should be rejected.

Appeal to NRT saves the classical account of predicative uses of comparative adjectives. For NRT straightforwardly predicts the readings we find in (31) and (32). According to NRT, each nominal co-occurs with a domain index. That is, according to the theory, the following are rough guides to the relevant aspects of the logical forms of (32) and (33):

- (34) Smith is a remarkable  $\langle$  violinist,  $f(i)$   $\rangle$ .  
 (35) Fred built a large  $\langle$  snowman,  $f(i)$   $\rangle$ .

Consider Kamp’s example. In both of Kamp’s envisaged scenarios, context assigns to “ $f$ ” a function from locations to people who have played instruments at those locations, and “ $i$ ” is assigned the salient location. In the context of the dinner party, the value of “ $i$ ” is the location of the dinner party. In the context of the London Festival Hall, the value of “ $i$ ” is the stage at the London Festival Hall. (32) is true relative to the first of these contexts, because Smith is remarkable compared to the violinists who have played in the past at the location at which the dinner party occurs. That is, where  $f$  is a function from locations to the people who have played instruments at those

locations, and  $l$  is the location of the dinner party, Smith is remarkable compared to the members of the intersection of  $\{x: \text{violinist}(x)\}$  and  $\{y: y \in f(1)\}$ . (31) is false relative to the second of these contexts, because Smith is not remarkable compared to the violinists who have played in the past at the location of the London Festival Hall. The case in which Fred is a 7-year-old child is similar. Relative to the envisaged context, “ $i$ ” could be assigned, for example, Fred, and “ $f$ ” a function from people to the set of structures that have been built by people of that age. Relative to a context in which “ $i$ ” and “ $f$ ” are assigned these values, (33) has the desired interpretation, that Fred built a snowman that is large for a snowman built by a 7-year-old child.

Now, let us turn back to predicative occurrences of comparative adjectives. The worry about the classical analysis of predicative occurrences of comparative adjectives was that the nominal by itself does not determine the comparison class. So, claiming that (26)–(28) are elliptical for (29)–(31) does not explain the existence of a contextually provided comparison class. But according to NRT, each nominal is really of the form  $\langle N, f(i) \rangle$ . Combining NRT with the classical account of predicative occurrences of comparative adjectives results in logical forms for (26)–(28) roughly like:

- (36) That building is a small  $\langle \text{building}, f(i) \rangle$ .
- (37) That basketball-player is a short basketball-  $\langle \text{player}, f(i) \rangle$ .
- (38) That flea is a small  $\langle \text{animal}, f(i) \rangle$ .

As we have seen, given the semantics provided in the previous section, instances of the schema “ $\langle N, f(i) \rangle$ ” do, relative to a context, determine the entire contextually salient comparison class for a comparative adjective.<sup>11</sup>

It has been standard in the literature on comparative adjectives to maintain that, even in attributive readings of adjectives, the nominal complement of a comparative adjective does not by itself determine the comparison class. What we have seen in this section is that this conclusion results from an inadequate grasp of the true syntax and semantics of nouns. Once NRT is adopted, one mystery about the “extra” context-dependency associated with the determination of comparison classes vanishes. The “extra” context-dependency in question is simply due to unrecognized structure in the noun, the very same structure that accounts for the phenomenon of so-called “quantifier domain restriction”.

<sup>11</sup> I have not here discussed degree theoretic approaches, despite their evident promise in yielding a unified account of adjectives and comparatives (degree theoretic approaches are first discussed with rigor in Cresswell’s classic (1975), and have recently been given new life in Kennedy (1997)). However, such approaches, as yet, have yielded no satisfactory analysis of predicative uses of adjectives. For example, the semantics for such constructions given in Kennedy (1997, 123 ff.) is non-compositional; the degree provided by context simply appears in the semantic derivation at the level of the degree phrase (for a discussion of the non-compositionality of such rules, see Chapter 2, pp. 104–6). A correct compositional treatment of predicative uses of adjectives will, I suspect, attribute the provision of the degree to the elided nominal in such constructions.

### III Mass Expressions

We have seen that NRT explains two apparently very different effects of context on linguistic interpretation. In this section, I discuss yet another apparently very different effect of context on linguistic interpretation that is explained by NRT. Consider the sentence:

(39) That puddle is water.

Suppose the puddle in question consists of muddy water. *Prima facie*, relative to certain contexts, (39) is true. However, *prima facie*, relative to other contexts, (39) is false. For example, suppose we are attending a conference of companies that market bottled water. Relative to such a context, it might be false that the stuff in the puddle counts as water. Indeed, relative to such a context, nothing less pure than the least pure bottled water might legitimately count as water. Finally, even more drastically, consider a context in which chemists are discussing the molecular structure of water. Relative to such a context, even the stuff in Evian bottles might not count as water.

Let us consider the denotation of “that puddle” to be fixed across contexts. One natural reaction to the *prima-facie* truth of (39) relative to some contexts, and the *prima-facie* falsity of (39) relative to other contexts, is to argue that, fixing the reference of the demonstrative expression “that puddle”, the truth of (39) does not vary from context to context. Relative to contexts in which “that puddle” refers to the relevant quantity of matter, (39) is either always false or always true. The impression otherwise is due to pragmatics, and not to semantics. Call this the *pragmatic account*.

One worry I have with the pragmatic account is that I do not see even the outlines of how a pragmatic account of these facts would proceed. For example, suppose that we take the view that each member of the relevant class of occurrences of (39) expresses a false proposition. I do not see how to provide a systematic, compelling Gricean derivation of the true proposition communicated from the false proposition expressed. Of course, Kent Bach and others have provided influential arguments that pragmatics also involves adding propositional constituents to the semantic content of an expression relative to a context, a process Bach calls “*Implicature*” (e.g., Bach (1994)). But even setting aside my worries about Bach’s notion of *implicature*, it is not clear to me how to provide an account based on *implicature* in defense of the pragmatic account of the above sort of context-dependency. In short, I worry that the pragmatic account puts too great of a strain on existing theories of pragmatics.

Of course, the countervailing worry is that incorporating this sort of contextual phenomenon into the semantics places too great of a strain on existing syntactic and semantic theories. There are two different ways to pose this worry. First, one might worry that incorporating the phenomenon into the semantics compromises the systematicity of semantic explanation in some deep fashion. Secondly, one might



worry that there is no independent evidence for the resources needed to treat this sort of context-dependence. If so, then postulating the mechanisms needed to treat this sort of context-dependence in the semantics may seem *ad hoc*.

The pragmatic account would be a fruitful avenue to explore if either of the above worries were legitimate. But, as it turns out, these worries are not legitimate. The resources needed to treat this sort of context-dependency are already in place. For NRT provides a smooth explanation for this sort of effect of context on semantic interpretation. NRT is independently motivated, and does not compromise the systematicity of semantic explanation.

According to NRT, each common noun co-occurs with a domain index. Now, as is well known, each count noun can be “transformed” into a mass expression. For example, “sailor” and “chicken” have mass-occurrences, as in (40a–b):

- (40) a. John had sailor for dinner. (John is a cannibal)  
 b. Hannah had chicken for dinner.

There is no reason to think that in using an expression that typically has count occurrences as a mass expression, one thereby drops the domain index. That is, there is no *prima facie* reason to think that, in (40a–b), the mass expressions do not co-occur with domain indices.

Furthermore, examination of mass quantification shows that there is just as much justification for the claim that each mass expression comes with a domain index as there is in the case of count quantification. For example, suppose that Pastor Hannah is concerned about the fact that someone has been drinking the holy water in her church on warm summer days. In a discussion with her John confesses:

- (41) I drank a little water last week.

What John expresses is the proposition that John drank a little of the church’s holy water the week before the utterance was made. That is, (41) is only true if John drank a little holy water, and not if he just drank a little unholy water.

We should expect any account of count quantifier domain restriction to generalize straightforwardly to the case of mass quantifier domain restriction. NRT, of course, does exactly this. With Helen Morris Cartwright, let us take an occurrence of a mass expression, relative to a circumstance of evaluation, to denote a set of quantities. According to NRT, the rough logical form of (41) is:

- (42) I drank a little  $\langle$  water,  $f(i)$   $\rangle$  last week.

In the context at hand, we may suppose “ $f$ ” to be assigned a function that takes a location to a set of quantities. Furthermore, “ $i$ ” is assigned Hannah’s church. Relative to this context, “ $f$ ” yields the set of quantities of holy liquids in the church. This set is then intersected with the set of all quantities of water, yielding the desired interpretation.

Furthermore, there is evidence that the object variable in structures such as (41) can be bound, as in:

(43) In every church, the pastor drinks a little water during the weekly ceremony.

(43) can express the proposition that in every church  $c$ , the pastor of  $c$  drinks a little of the holy water of  $c$  during the weekly ceremony. Thus, as one might expect, there is the same sort of evidence for mass quantifier domain variables in logical forms as there is for count quantifier domain variables.

According to NRT, then, (44) is a rough guide to the relevant aspects of the logical form of (39):

(44) That puddle is  $\langle \text{water}, f(i) \rangle$ .

Relative to a circumstance of evaluation, “water”, in (39), denotes the set of all quantities of water. Relative to a context, “ $f$ ” is assigned a function that takes an object of some kind and yields a set of quantities of matter, one appropriate to intersect with the denotation of “water”. Relative to a context, “ $i$ ” is assigned an object that serves as an input to the denotation of “ $f$ ”.<sup>12</sup>

If NRT is true, then the sort of context-dependence at issue in examples such as (39) is easily explicable. Consider the context of the bottled water conference. In this case, context assigns to “ $f$ ” a function that takes a location and yields a set of quantities, and context assigns to “ $i$ ” the bottled water conference. The result of applying the value of “ $f$ ” to the bottled water conference will be, for example, the set of quantities of liquids that are sufficiently pure to be sold in bottles in supermarkets in the United States. The denotation of “ $\langle \text{water}, f(i) \rangle$ ” is then the result of intersecting the denotation of “water” with this set of quantities. Similarly, the context in which chemists are discussing the molecular structure of water may be one in which the set of contextually provided quantities contains all and only those quantities  $q$  such that  $q$  is constituted by molecules of the same type. Finally, the context in which (39) counts as true is one in which the domain is all quantities whatsoever, so the denotation of “water” is unrestricted.

So NRT provides a straightforward explanation of the sort of context-dependence in examples such as (39). There is thus evidence from a wide variety of constructions for the truth of NRT. Furthermore, given NRT, many apparently distinct effects of extra-linguistic content on what is asserted are traceable to the same source.

<sup>12</sup> One might worry that the value of “ $f(i)$ ”, relative to a context, is not a plausible mass denotation, since it is not closed under sums (a property, according to Quine (1960, 91), that is constitutive of mass-denotations). However, I doubt that mass denotations are closed under sums. One example, due to George Boolos, is “dust”. Not every sum of quantities of dust is dust. Another example, due to Zsófia Zvolensky, is “liquid”. There might be two liquids that, when combined, turn into a gas (however, in this latter case, one might respond by denying that a sum is a mixture).

It is worthwhile noting a few consequences of this analysis that are of general philosophical interest. According to the analysis we have given, and supported by syntactic and semantic evidence from a wide variety of constructions, the reading on which (39) is true is one according to which “water” occurs unrestricted. That means that the denotation of an unrestricted use of “water” includes quantities that are very far from pure  $H_2O$ . Exactly parallel considerations govern all other mass expressions that are so-called “natural kind terms”, such as “gold”.

In the theory of reference, it is often assumed that “water” denotes  $H_2O$ . The first consequence of our analysis is that this assumption is incorrect. The denotation of “ $H_2O$ ”, namely the set of all quantities of  $H_2O$ , is but a small proper subset of the denotation of “water”. Though in some contexts, “Water is  $H_2O$ ” expresses a truth, this is not because the literal meaning of “water” is the same as the literal meaning of “ $H_2O$ ”. In some contexts, “Every man is a judo-expert” expresses a truth, but we should not infer from this that the literal meaning of “man” is the same as the literal meaning of “judo-expert”. Neither inference is valid.

A second consequence of note that can be drawn from the foregoing concerns expressions such as “literally speaking”. Perhaps there are true uses of:

(45) Literally speaking, only pure  $H_2O$  is water.

But if there are, then what it shows is that the function of “literally speaking” is not to restrain the interpretation of the words used to their literal meanings. The literal meaning of “water” determines an extension that includes quantities that contain molecules distinct from  $H_2O$ . The function of “literally speaking” in (45) is to restrict the domain of “water” to a small sub-domain of its extension. In a sense, then, literally speaking is quite far from speaking literally.

One worry one might have about the account I have given is that it makes the extension of “water” implausibly large. For example, one might worry that any substance whatsoever, in some context, counts as water. This worry is misplaced, if one accepts that, for every noun, there is at least one context in which it can be used with its domain maximally wide. In the case of mass expressions, what this means is that, for each mass noun, it is possible for there to be a context in which the domain for the mass noun is the set of all quantities whatsoever. Together with the fact that there is no true use of:

(46) Any quantity of any old substance whatsoever is water.

we may conclude that the literal meaning of “water” does not determine the set of all quantities whatsoever.<sup>13</sup>

<sup>13</sup> Thanks to Tim Williamson for discussion here.

Of course, I have not seriously addressed the question of what fixes the reference of a term such as “water”. But the results we have seen should serve to temper any inclinations one may have to draw exaggerated consequences from the view that terms such as “water” have “hidden scientific essences”. No doubt, for a quantity of liquid to count as a quantity of water, it must contain a certain portion of  $H_2O$  molecules. But a quantity of blood may contain a greater percentage of  $H_2O$  molecules than some quantity of water, without thereby counting as a quantity of water. Spelling out what fixes the reference of a term like “water” is a difficult matter. But, given that the diluted stuff in lakes is water, such a story may very well centrally involve the sort of information available to ordinary speakers competent with the term “water”, such as “falls from the sky in the form of rain”. This is evidence, albeit from an unlikely source, for an externalist, description theoretic account of the meanings of words such as “water” and “gold”.

As we have seen, (non-contextually restricted uses of) terms such as “water” and “gold” do not denote pure chemical kinds. This result supports a view recently advocated by Mark Johnston (1997). According to Johnston, mass terms such as “water” do not denote chemical kinds such as  $H_2O$ . Rather, they denote what Johnston calls “manifest kinds”. The manifest kind denoted by “water” is not identical to  $H_2O$ . Instead, it is constituted by  $H_2O$ .

For a quantity to count as a quantity of water, perhaps it must contain some  $H_2O$ . But, as we have seen, it may contain less  $H_2O$  than a quantity that is clearly not a quantity of water. These facts provide fairly decisive evidence for the thesis that water is not identical to  $H_2O$ .

Furthermore, our discussion provides a decisive rebuttal to one natural response to Johnston’s arguments that “water” and similar mass terms denote “manifest” kinds. According to this response, each such term is ambiguous between a “scientific” use, in which it denotes a chemical kind, and a “manifest” use, where it functions as Johnston describes. However, as we have seen, uses of “water” that denote a set of quantities, each of which is pure  $H_2O$ , do not involve a separate lexical item. Rather, these uses involve the very same lexical item as more ordinary uses of water. However, in these specialized uses, the denotation of that lexical item is contextually restricted. As far as literal meaning is concerned, then, terms such as “water”, “gold”, “copper”, and the like uniformly do not denote chemical kinds.<sup>14</sup>

<sup>14</sup> I should note that discussions with Richard Boyd have made me somewhat uncomfortable with the dichotomy suggested by Johnston’s terminology of “manifest” vs “chemical” kind. The fact that (non-contextually restricted uses of) “water” and “gold” fail to denote chemical kinds does not entail that (such uses of) these terms do not denote natural kinds. After all, chemical kinds are not the only natural kinds. It is unclear to me that explanation in, say, evolutionary biology, can do without appeal to the kinds denoted by “water” and “gold”, despite the fact that they are not chemical kinds.

## Conclusion

There are other constructions the analysis of which is aided by the adoption of NRT. For example, like indefinite descriptions, definite descriptions exhibit quantificational variability effects. That is, among the interpretations of (47 a–b) are the ones specified in (48 a–b):

- (47) a. The customer is always right.  
 b. Usually, the sailor stops, but the marine goes on.
- (48) a. For all  $x$ , if customer ( $x$ ), then right ( $x$ ).  
 b. For most  $x, y$ , if sailor ( $x$ ) and marine ( $y$ ), then  $x$  stops and  $y$  goes on.<sup>15</sup>

Such effects are surprising, on the supposition that definite descriptions are quantifiers. Indeed, Delia Graff (2001, section vii) has recently exploited the fact that definite descriptions are subject to quantificational variability effects to argue for the thesis that definite descriptions are not quantifier expressions, but predicates. For if definite descriptions were predicates, one could explain the data in (47) and (48), since the adverb of quantification would then bind a free variable in the predicate expression.

However, NRT reconciles the quantificational treatment of definite descriptions with the existence of quantificational variability effects. For, if NRT is correct, each noun co-occurs with a domain index. That is, according to NRT, the structure of (47 a–b) is as in:

- (49) a. The  $\langle$  customer,  $f(i)$   $\rangle$  is always right.  
 b. Usually, the  $\langle$  sailor,  $f(i)$   $\rangle$  stops, but the  $\langle$  marine,  $f(j)$   $\rangle$  goes on.

Given these representations, NRT smoothly predicts quantificational variability effects. In (47a), the adverb of quantification “always” raises and binds the variable “ $i$ ” in  $\langle$  customer,  $f(i)$   $\rangle$ . We may suppose then that context supplied “ $f$ ” with a function from situations to sets such that the intersection with that set of the denotation, relative to a circumstance of evaluation, of “customer”, yields a set with one member. Similarly, in (47b), the adverb of quantification “usually” unselectively binds the variables “ $i$ ” and “ $j$ ”, yielding the desired interpretation.<sup>16</sup> Thus, NRT allows us to explain such effects, without abandoning the thesis that definite descriptions are quantifiers.

However, my point in this chapter is not simply to emphasize the virtues of the Nominal Restriction Theory. It appears that the effects of extra-linguistic context on the determination of what is said by the use of a sentence are too diverse and

<sup>15</sup> I am here ignoring the complexities involved in the proportion problem (see for discussion, e.g., Heim (1982); Reinhart (1986)).

<sup>16</sup> For more on adverbs of quantification and unselective binding, cf. Lewis (1975); Heim (1982).

varied to be susceptible of systematization. Appeals to extra-linguistic context are consequently ubiquitous in the work of those who seek to undermine the thesis that linguistic interpretation is largely systematic and rule-governed. However, I hope to have shown by example how what appear to be very different effects of context on the determination of what is said can be due to the same source. If, as I believe to be the case, there are only a small number of ways in which extra-linguistic context affects what is said by a use of a sentence, perhaps it is not so clear that extra-linguistic context poses a threat to the systematicity of linguistic interpretation.

## References

- BACH, K. (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BARWISE, J., and COOPER, R. (1981), 'Generalized Quantifiers and Natural Language', *Linguistics and Philosophy*, 4: 159–219.
- CHIERCHIA, GENNARO (1993), 'Questions with Quantifiers', *Natural Language Semantics*, 1: 181–234.
- CRESSWELL, MAX (1975), 'The Semantics of Degree', in Barbara Partee (ed.), *Montague Grammar* (New York: Academic Press), 261–92.
- ENGDAHL, ELISABET (1986), *Constituent Questions* (Dordrecht: Reidel).
- FARKAS, DONKA (1997), 'Evaluation Indices and Scope', in A. Szabdczi (ed.), *Ways of Scope Taking* (Dordrecht: Kluwer), 183–215.
- GRAFF, DELIA (2001), 'Descriptions as Predicates', *Philosophical Studies*, 102/1: 1–42.
- GROENENDIJK, J., and STOKHOF, M. (1984), 'Studies on the Semantics of Questions and the Pragmatics of Answers' (Amsterdam doctoral diss.).
- HEIM, IRENE (1982), 'The Semantics of Definite and Indefinite Noun Phrases' (doctoral diss., University of Massachusetts at Amherst).
- HEIM, IRENE, and KRATZER, ANGELIKA (1998), *Semantics in Generative Grammar* (Oxford: Blackwell).
- JOHNSTON, MARK (1997), 'Manifest Kinds', *Journal of Philosophy*, 94/11: 564–83.
- KAMP, HANS (1975), 'Two Theories of Adjectives', in E. Keenan (ed.), *Formal Semantics of Natural Language* (Cambridge: Cambridge University Press), 23–55.
- KENNEDY, C. (1997), *Projecting the Adjective: The Syntax and Semantics of Gradability and Companion*. diss. University of California (Santa Cruz, Calif.: SLUG pubs.).
- LARSON, RICHARD (1998), 'Events and Modification in Nominals', in *Semantics and Linguistic Theory*, VIII (Ithaca, N. Y.: CLC), 145–68.
- LARSON, RICHARD, and SEGAL, GABRIEL (1995), *Knowledge of Meaning* (Cambridge, Mass.: MIT Press).
- LEWIS, DAVID (1975), 'Adverbs of Quantification', in E. Keenan (ed.), *Formal Semantics of Natural Language* (Cambridge: Cambridge University Press).
- PARSONS, TERENCE (1972), 'Some Problems Concerning the Logic of Grammatical Modifiers', in E. Harman and D. Davidson (eds.), *Semantics of Natural Language* (Dordrecht: Reidel), 127–41.

- QUINE, WILLARD VAN ORMAN (1960), *Word and Object* (Cambridge, Mass.: MIT Press).
- REINHART, TANYA (1986), 'On the Interpretation of Donkey Sentences', in E. Traugott *et al.* (eds.), *On Conditionals* (Cambridge: Cambridge University Press).
- SIEGEL, MUFFY (1975), 'Capturing the Russian Adjective', in Barbara Partee (ed.), *Montague Grammar* (New York: Academic Press), 293–309.
- WESTERSTÅHL, DAG (1989), 'Quantifiers in Formal and Natural Languages', in D. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic, IV* (Dordrecht: Kluwer), 1–131.

# 4

## Semantics, Pragmatics, and the Role of Semantic Content

(with Jeffrey C. King)

Followers of Wittgenstein allegedly once held that a meaningful claim to know that  $p$  could only be made if there was some doubt about the truth of  $p$ . The correct response to this thesis involved appealing to the distinction between the *semantic content* of a sentence and features attaching (merely) to its *use*. It is inappropriate to assert a knowledge-claim unless someone in the audience has doubt about what the speaker claims to know. But this fact has nothing to do with the semantic content of knowledge-ascriptions; it is entirely explicable by appeal to pragmatic facts about felicitous assertion (that is, a kind of use of a sentence).

According to the contextualist about knowledge, the (propositional) semantic content of knowledge-claims is sensitive to context. In a context in which skeptical possibilities are sufficiently salient, the word “know” expresses a relation that holds between persons and a highly restricted range of propositions. In a context in which skeptical possibilities are not salient, the word “know” expresses a different relation, one that a person can bear to a proposition even if she is in a fairly weak epistemic position with respect to it. In support of her position, the contextualist often points

We have given this paper at the Kentucky Language Association, the Universidad Nacional Autónoma de México, the Oxford Philosophical Society, and it was also discussed at the Bay Area Philosophy of Language Reading Group; thanks to participants at all of these events. Herman Cappelen, Robyn Carston, Maite Ezcurdia, Richard Heck, Ernie Lepore, and Rich Thomason deserve particularly special thanks. Kent Bach, Michaela Ippolito, Matthew Stone, Zoltán Gendler Szabó, and Brian Weatherston also helped with valuable suggestions. Thanks also to two anonymous referees for Oxford University Press.



to the undisputed fact that speakers' willingness to make knowledge-claims varies with context. Those who reject contextualism about knowledge typically try to give non-semantic accounts of variations in speaker hesitancy about knowledge-claims. In short, dissenters from contextualism try to argue that the facts that support contextualism are really facts about the *use* of knowledge-ascriptions, rather than their *semantic contents*.

In ethics, one version of internalism about reasons holds that someone understands a sentence containing genuinely normative vocabulary only if they are motivated in a certain way. That is, an internalist about moral reasons holds that being motivated is an essential part of the grasp of the semantic content of moral sentences. An internalist about reasons motivates her position in part by appealing to the fact that it is odd to utter an ethical sentence unless one has the relevant motivation. An externalist about reasons, by contrast, rejects the internalist thesis that being motivated is part of grasping the semantic content of moral sentences. An externalist seeks to explain the evidence about motivation by attributing it to facts (merely) about the proper *use* of ethical sentences, rather than the *semantic content* thereof.

These examples form but a representative few. Appealing to a distinction between the semantic content of a sentence and what a use of it only pragmatically conveys is simply a standard maneuver in debates across a wide range of disciplines in philosophy. The importance of the tactic is due to the nature of many philosophical claims. Typically, a philosopher who is discussing a certain discourse, whether it concerns knowledge or the good, makes claims about its content. Sometimes, there are features of the use of a word in that discourse that, if explained by reference to the content of the word, would threaten that philosopher's claims. In such a case, the philosopher who continues to advocate her claims has one of two options. First, she can reject the thesis that the ordinary word that purportedly expresses that content in fact expresses it (typically, this results in an error theory about ordinary discourse). Secondly, she can argue that the features in question are not explained by reference to the semantic content of that word, but are merely pragmatic facts about its use. In short, the distinction between *semantics* and *pragmatics* is fundamental to philosophical theorizing, because much philosophical theorizing takes the form of claims about the content of philosophically central discourse.

So, in a number of debates across, for example, metaphysics, epistemology, and ethics, one theorist's semantic content is another theorist's merely pragmatic effect. But a complicating factor in these debates is the lack of a clear and accepted criterion among philosophers of language and linguists for what counts as semantic versus what counts as pragmatic. That is, among philosophers of language, there is no stable agreement on the semantics–pragmatics distinction. Furthermore, even among those who agree on terminology, there is disagreement about the *scope* of semantic content. That is, it is a subject of much current debate how much of what is intuitively

communicated is constituted by semantic content. Many philosophers of language have in recent years argued that sentences have only minimal, even non-propositional, semantic contents, much of their natural interpretation in a context being due to non-semantic effects (e.g. Bach (1994)). Others in the philosophy of language and cognitive science community have gone further, to reject altogether any theoretical role for a semantic notion of what is said by a sentence. According to all of these theorists, instead of “truth-conditional semantics”, we should be speaking, in the words of François Recanati, of “truth-conditional pragmatics” (Recanati (1993, 232 ff.)).

Our first aim in this chapter is to provide a stable characterization of the distinction between semantics and pragmatics, one that theorists across a range of disciplines can exploit. This characterization will allow us to gain greater clarity about the debates concerning the scope of semantic content. Our second aim in this chapter is to argue that much more counts as genuinely semantic than skeptics about the scope of semantic content have maintained.

## I Three Conceptions of the Semantics–Pragmatics Distinction

Above, following tradition, we spoke of the distinction between semantics and pragmatics as a distinction between what words mean (semantics), on the one hand, and the *use* speakers make of words (pragmatics). However, this characterization is imprecise and unhelpful. For example, the case of indexical expressions, such as “I” and “today”, shows that one word can have different denotations on different occasions of use. Nevertheless, each indexical word type has a univocal conventional meaning, that is, a meaning that does not vary from context to context. Since the referential content of a use of “I” relative to a context of use seems to be a function of the conventional meaning of the word type “I”, it seems incorrect to relegate all facts about its content in a context to what speakers use “I” to mean, rather than what the word means. But the initial vague characterization of the semantics–pragmatics distinction does not talk of word meaning relative to a context, and so gives us no handle on these facts. More precision is required.

Following Richard Heck (2001), let us call the conventional meaning of a lexical item *e* in a language *L* its *standing meaning*.<sup>1</sup> Some terms, such as the number determiners

<sup>1</sup> David Kaplan’s (1989) technical reconstruction of standing meaning is what he calls character; John Perry’s classic terminology is “role”. Since we wish here to remain neutral about certain (albeit differing) theoretical commitments that, strictly speaking, Kaplan and Perry’s terminology possesses, we adopt Heck’s terminology for now.

“two” and “three”, or proper names such as “Bill Clinton” and “George Bush”, seem to have a “stable” standing meaning, in the sense that users using such terms correctly (that is, in accord with their standard meanings), will always refer to the same object or property. Other terms, such as “I”, “here”, or “this”, have “unstable” standing meanings, such that, in different contexts, consistently with these standing meanings they can be used to refer to different objects. For example, the conventional meaning of “I” in English does not vary across contexts; in every context, when used in accord with its standard meaning, the meaning of “I” is (roughly) the same as “the speaker in the context”. But in a context in which Bill Clinton is the speaker, Bill Clinton uses “I” to refer to himself, when using “I” in accord with its standing meaning. In contrast, Gray Davis uses “I” in accord with its standing meaning to refer to himself, that is, Gray Davis. Loosely following recent terminology of John Perry, let us call the object or property that a user refers to in a context when successfully using a lexical item *e* (of a language *L*) in accord with its standing meaning in *L* the *referential content* of that term relative to that context.<sup>2</sup>

There is a bit of a terminological morass surrounding our distinction between “standing meaning” and “referential content”. The term “referential” is most appropriate when used to discuss the content of *referential expressions*, expressions that can be used to refer, as opposed to merely denote, in the sense familiar from the work of Bertrand Russell. Such expressions, following tradition, are assumed to be singular terms, paradigmatically proper names such as “Bill Clinton” and indexicals and demonstratives such as “I” and “this”. However, like Perry (2001, 79), we intend the expression “referential content” to have wider application than merely to referential expressions in contexts. Our distinction between standing meaning and referential content is meant to correlate with the distinction between, on the one hand, the context-invariant standing meaning of a term, and, on the other, the object, property, or function that the term has as its content in a context, which is potentially distinct from its standing meaning. As a consequence, the expression “referential content” may need to have wider application than merely to the class of referential expressions. For there may be context-sensitive elements in the syntax of a sentence that are not referential expressions, as traditionally conceived (e.g. Heal (1997)). If there are, we would need vocabulary to mark the distinction between standing meaning and content in a context for them as well. So, like Perry, we will speak of the referential content of an expression in a context in this broader sense.

The distinction between the non-relative notion of the standing meaning of a term and the relative notion of the referential content of a term relative to a context provides

<sup>2</sup> See ch. 5 of Perry (2001). We say “loosely”, because Perry takes referential content to be a property of *utterances* of expressions. In contrast, we take referential content to be a property of an expression relative to a context.

two general strategies for forging a distinction between semantics and pragmatics, and a third intermediate between them. We discuss each in turn.

According to the first strategy for making a distinction between semantics and pragmatics, the semantic interpretation of a complex expression *e* is the result of composing the standing meanings of the lexical items in *e* in accordance with the semantic composition rules corresponding to the syntactic structure of *e*. So, on this view, semantic properties are only properties of expression types; any property that an expression type has only relative to a context (or any property possessed only by expression tokens) is not semantic. The semantic content of the sentence “I am tired” is the result of combining the standing meaning of “I”, the standing meaning of the copula (together perhaps with the present tense), and the standing meaning of “tired”.<sup>3</sup>

Of course, relative to a particular context of use, speakers can use “I” in accord with its standing meaning to refer to a particular person, namely the speaker of that context. But on this first conception of semantics, semantic properties do not accrue to expressions in context. The fact that in a particular context of use, “I” has Bill Clinton as its referential content is a fact about the pragmatics of “I”, how it is used in a given context.

Propositions are the ultimate bearers of truth-values. This fact about the nature of propositions entails that the proposition expressed by “I am a son of a president” varies from context to context. Relative to a context in which Bill Clinton is the speaker, the proposition expressed by “I am a son of a president” is false, whereas relative to a context in which George Bush is the speaker, the proposition expressed is true. However, on this first way of thinking of semantics, no semantic value of an expression varies from context to context. Therefore, on this first way of distinguishing semantics from pragmatics, the *semantic* interpretation of “I am a son of a president” is not a proposition.

This conception of the semantics–pragmatics distinction is familiar from the work of Richard Montague. Montague suggests that the difference between semantics and pragmatics is that semantic values are not relativized to anything (or rather only to models or interpretations of the language), whereas pragmatic values are assigned relative to a context of use (as well as a model). Thus he writes: “It seemed to me desirable that pragmatics should at least initially follow the lead of semantics, which is primarily concerned with the notion of truth (in a model, or under an interpretation), and hence concern itself also with truth—but with respect to not only an interpretation but also a context of use” (Montague (1974, 96)). Few linguists or philosophers now conceive of the distinction between semantics and pragmatics in this first way (though, as we emphasize in the next section, the reasons that justify abandoning it are subtle). According to a more contemporary conception of the semantics–pragmatics distinction, there are two levels of semantic value. The first

<sup>3</sup> Of course, the syntactic structure of “I am tired” is certainly considerably more complex than meets the eye, and presumably contains many non-obvious elements.

is the non-relativized notion of the standing meaning of an expression. The second is the relativized notion of the referential content of an expression in a context.<sup>4</sup> On the most influential way of explaining the relation between these two levels of semantic value, the standing meaning of a lexical item determines a *function* that determines its referential content, given a context. Following the influential terminology of Kaplan (1989), the widely used name for this function is *character*, in what follows, when speaking of this formal explication of standing meaning, we adopt Kaplan's terminology. Most importantly, a complex expression relative to a context *c* has a referential content that is the result of combining the referential contents of its constituent terms relative to the context *c* in accord with the semantic composition rules corresponding to the syntactic structure of that expression. The result of this latter process is a genuine level of semantic value, which we shall call the *semantic content* of that complex expression relative to the context *c*.<sup>5</sup> The semantic content of a lexical item relative to a context *c* is, on this view, its referential content in that context.

On this second conception of the semantics–pragmatics distinction, the sentence “I am a son of a president”, relative to a context, has a propositional semantic content. If George is the speaker, and the time is *t*, the propositional content is that function that takes a possible situation *s* to the true if and only if George is a son of a president at *t* at *s*. On this conception of the semantics–pragmatics distinction, the pragmatic content is what the speaker communicates over and above the semantic content of the sentence he uttered.

So, for example, suppose George is asked whether he thinks he is going to be rich when he is older, to which he responds “I am a son of a president”. The semantic content of George's sentence, relative to that context, is that he is a son of a president, which is the proposition that results from combining the referential contents, relative to that context, of the parts of the sentence. But by expressing this semantic content, he intends to flout Grice's maxim of relevance, thereby communicating the quite different proposition that it is quite likely that he will be rich, given his political

<sup>4</sup> Some philosophers of language have advocated semantic theories that embrace more than two levels of semantic content. For example, Nathan Salmon (1986) holds that an adequate account of tensed discourse requires at least three levels of semantic value, and David Braun (1996) argues that an adequate treatment of demonstratives requires at least three levels of semantic value. However, for the sake of simplicity, we ignore in what follows the possibility that there are more than two levels of semantic value.

<sup>5</sup> There are basically two different styles of semantic theory. For the sake of simplicity, we take the content of an expression in a context to be a function from possible worlds to the appropriate type of extension; functions from possible worlds to truth-values we will call *propositions*. On a structured propositions semantics, however, semantic interpretation proceeds in two steps. First, an algorithm assigns structured propositions to sentences in contexts. Secondly, a definition of truth assigns truth-conditions to structured propositions. In what follows, we generally speak in the first style of semantics, but translate into the second in some instances.

connections. This proposition is communicated not by being semantically expressed by some sentence, but rather as a post-semantic result of the expression of a semantic value. Hence, it is part of the pragmatic content of the speech act, but not what is semantically expressed by the sentence in that context.

On this conception of the semantics–pragmatics distinction, the semantic contents of some terms depend upon non-linguistic context. Those who use “pragmatic” to refer to such contextual effects on linguistic interpretation may speak, somewhat paradoxically, of pragmatic effects on semantic content. But what this means is simply that some lexical items have different contents in different contexts. Non-linguistic facts about the context of use are relevant for fixing the referential content of some lexical items, such as pronouns and unpronounced free variables. But the nature of the lexical item dictates what non-linguistic facts are relevant, and constrains the nature of its referential content in a context. For example, if the lexical item is the English pronoun “she”, its standing meaning dictates that only certain sorts of intentions are relevant for fixing its referential content in a context, and constrains that content to be, for example, a salient female human (reference to boats, countries, etc., to one side). That is, speaker intentions are relevant to fixing the referential content of a lexical item in a context only when they are determined to be so by the standing meaning of a lexical item. So, the role played by speaker intentions in semantics remains significantly constrained, even on this conception of the semantics–pragmatics distinction, by the standing meanings of lexical items.

Employing this version of the semantics–pragmatics distinction, we can distinguish between two ways in which context determines what is communicated. The first way context may determine what is communicated is by affecting the semantic content, via resolution of the referential content of context-sensitive elements in the sentence uttered. This roughly corresponds to what Stanley and Szabó (pp. 79–80) and Perry (2001, 42 ff.) call the *semantic* role of context.<sup>6</sup> The second way is that context plays a role in determining what is communicated by the linguistic act over and above its semantic content. This is the genuinely pragmatic role of context (p. 80).<sup>7</sup>

<sup>6</sup> We say “roughly corresponds”, since Stanley and Szabó make a point of not adopting the semantics–pragmatics distinction under discussion here (pp. 79–80). Rather, according to them, the semantic interpretation of an utterance is the proposition it expresses, whether or not all elements in the proposition expressed by the utterance are traceable to elements in the sentence uttered. So what Perry (2001, 44 ff.) calls the “content supplemental” role of context is, for Stanley and Szabó, a semantic role of context. We find this terminology disturbing, since it is unclear whether such a notion of the semantics–pragmatics distinction takes semantic content to be a property of expressions in contexts or utterances.

<sup>7</sup> Stanley and Szabó also discuss what they call the “grammatical role” of context in determining what was uttered, which corresponds directly to what Perry (2001, 40 ff.) calls the “pre-semantic” role of context. This is the role context plays in resolving ambiguity. Both Perry and Stanley and Szabó are

So, on this second conception of semantics, extra-linguistic context plays a role in determining the semantic content of certain expressions in context. Indeed, for many context-dependent expressions, certain kinds of intentions of speakers will likely be relevant for determining the semantic content of some expressions relative to a context. But this does not undermine the fact that these values are the semantic contents of the relevant expressions, given this semantics–pragmatics distinction. In accord with this, one can define two senses of “pragmatic effect”, which we will henceforth call “weak pragmatic effects” and “strong pragmatic effects”. A weak pragmatic effect on what is communicated by an utterance is a case in which context (including speaker intentions) determines interpretation of a lexical item in accord with the standing meaning of that lexical item.<sup>8</sup> A strong pragmatic effect on what is communicated is a contextual effect on what is communicated that is not merely pragmatic in the weak sense.<sup>9</sup>

On the first conception of the semantics–pragmatics distinction we discussed, there are no pragmatic effects of context on semantic content, whether weak or strong. If there were, for example, a weak pragmatic effect on the semantic content of some expression, then its semantic content would have to be relativized to contexts, to incorporate this sensitivity. But there are no relativized semantic values on this first conception. On the second conception of the semantics–pragmatics distinction, the semantic content of a complex expression in a context is a function of, and only of, the referential contents of its constituents in that context, together with context-independent composition rules. A strong pragmatic effect on what is communicated is one in which context affects what is communicated, but not by affecting the referential contents of any lexical item in a sentence. So, by definition, there are no strong pragmatic effects on semantic content on the second conception of the semantics–pragmatics distinction.<sup>10</sup> However, unlike the first conception, all weak pragmatic effects of context are cases in which context has a semantic

too quick to lump these roles of context together. For the sense of “determine” here is quite distinct from the sense of “determine” we have exploited. Facts about context determine the semantic contents of contextually sensitive items, and the implicatures of the speech act. In contrast, facts about context do not determine which of two disambiguations a given utterance expresses; it is rather that hearers draw on context to figure out which unambiguous utterance was expressed.

<sup>8</sup> The purpose of the last proviso is to rule out deferred reference, cases such as “The ham sandwich is getting irritated” (said by a waiter), as a weak pragmatic effect, since such cases are not ones in which context works to determine interpretation in accord with the standing meaning.

<sup>9</sup> Resolving ambiguity is not a sense in which context “determines” what is communicated, and so counts as neither a strong nor a weak pragmatic effect on what is communicated.

<sup>10</sup> The Context Thesis of Zoltán Gendler Szabó is the principle that “The content of an expression depends on context only in so far as the contents of its constituents do” (Szabó (2001, 122)). Szabó’s Context Thesis, on the second conception, is a definitional truth about semantic content.

role. That is, all weak pragmatic effects are cases in which context affects semantic content.<sup>11</sup>

Some theorists, torn between the restrictiveness of the first conception and the perceived permissiveness of the second, choose to split the difference between these ways of approaching the semantics–pragmatics distinction. One way of adopting a semantics–pragmatics distinction intermediate between the first and second conceptions is to distinguish between different kinds of weak pragmatic effects. One can then use the resulting distinction to count some weak pragmatic effects as affecting semantic content, and others as not having effects on semantic content. Here is one way to motivate this position. As John Perry has emphasized, context-dependent expressions come in two classes.<sup>12</sup> First, there are what Perry calls “automatic” indexicals. Two examples of such indexicals, in English, are “I” and “tomorrow”. The referential content of these indexicals, in a context, is fixed independently of the beliefs or intentions of its user.<sup>13</sup> An occurrence of ‘tomorrow’ has, as its referential content, the day after the day it is used, independently of what the speaker intended it to refer to, or believed it referred to.<sup>14</sup> The vast majority of other context-sensitive expressions are what Perry calls “intentional”. In the case of, for example, a use of ‘that man’, it is the speaker’s intentions that help determine its referential content in a context.<sup>15</sup>

Given Perry’s distinction between automatic and intentional context-sensitive expressions, one might adopt the following distinction between semantics and pragmatics. The semantic content of a sentence in a context is a function of (and only of) the referential contents of the automatic indexicals in the sentence relative

<sup>11</sup> One might worry that the second conception rules out strong pragmatic effects on semantic content only in spirit. For example, consider a purported strong pragmatic effect on what is communicated by a use of a sentence *S*. One could imagine transforming this into a weak pragmatic effect, by making some expression in *S* into an indexical. By this retranslation scheme, one could transform strong pragmatic effects into weak pragmatic effects. However, this maneuver is unpersuasive; the basic reply is that one cannot just *stipulate* that a word is context-sensitive. Given that context-sensitive expressions are typically identifiable, there is a high burden of proof on someone who wishes to maintain that a non-obviously context-dependent expression is in fact context-dependent (pp. 64–5). Furthermore, there are tests to distinguish genuine indexical expressions from non-indexicals; for example, indexicals are invariant in interpretation under verb–phrase ellipsis, see Stanley (2003).

<sup>12</sup> See ch. 4 of Perry (2001). See also Perry (1997), which is a precursor.

<sup>13</sup> Except that, as Perry (2001, 596) points out, the user must intend the indexicals to be used in accord with their standard meaning.

<sup>14</sup> A *locus classicus* of this point, at least for “I”, is Wettstein (1984).

<sup>15</sup> As Ernie Lepore has pointed out to us, some philosophers, such as Chris Gauker, reject the existence of intentional indexicals. These philosophers are thoroughgoing contextualists; they hold that non-mental features of the context determine the value of all contextual parameters relevant for determining what is said (see Gauker (1997) for a good exemplar of this tradition). For a contextualist, the third conception of the semantics–pragmatics distinction collapses into the second.



to that context, together with the standing meanings of all other lexical items in the sentence, together with context-independent composition rules. On this account, the semantic content of “I am tired” differs from context to context, since “I” is an automatic indexical. But the semantic content of “that man is tired” does not differ from context to context, since “that man” is an intentional context-sensitive expression, and so has its referential content in a context fixed by appeal to speaker intentions, which cannot, on this view, be relevant to the determination of semantic content.<sup>16</sup>

We have now discussed three different ways of drawing the distinction between semantics and pragmatics. We think the second of these three ways of drawing the distinction is most plausible. There are powerful phenomenological considerations in its favor. Consider a sentence relative to a context. The result of composing the referential contents of its parts leads to a quite natural entity to take as its semantic content in that context, one that seems to play a fundamental role in interpretation, as the object of speech acts and attitudes. For example, the result of combining the referential contents of the parts of the sentence “She is tired now”, relative to a context in which the speaker intends to refer to Hannah and *t* is the time of utterance, is the proposition that Hannah is tired at *t*. This seems to be the object of propositional attitudes and what is understood in a successful case of communication. So, the entity that results from combining the referential contents of the parts of a sentence relative to a context in accord with composition rules seems to be one that plays a recognizably central role both in theorizing about linguistic interpretation, and in giving an adequate account of the semantic contents of some linguistic constructions.

But the main reason for favoring this second way of drawing the distinction between semantics and pragmatics is that there are difficulties facing the first and third ways of drawing the distinction between semantics and pragmatics, the resolution of which leads to the adoption of our favored conception. We devote the next two sections to a detailed discussion of these difficulties.

## II Semantic Content as Context-independent Semantic Value

On the first way of drawing the semantics–pragmatics distinction, semantics is only concerned with assigning *unrelativized* semantic values to sentences. Accordingly, any

<sup>16</sup> This third conception of the semantics–pragmatics distinction is, for example, clearly operative in the work of Kent Bach; as he writes, “Contextual information in the narrow, semantic sense is limited to a short list of parameters associated with indexicals and tense, such as the identity of the speaker and hearer and the time of the utterance” (Bach (2002a, 285)).

value assigned to a sentence that is relativized to a context, or point of evaluation, etc. (other than a model for the language) falls within the domain of pragmatics. We know of three ways such a view might be implemented. In this section we take each in turn and provide criticisms of them.

The first way of implementing the proposal is inspired by the work of Richard Montague. One assigns an unrelativized semantic value to expressions of the language (again, this semantic value may be relativized to a model or “possible interpretation” for the language, but it is not relativized to anything like contexts or points of evaluation; we will suppress this qualification henceforth). This value is a function from “points of reference” to appropriate extensions. Though the need for “double indexing” was recognized after the work of Montague on these topics, one can derive an updated version of Montague’s approach by letting the points of reference be pairs of contexts and points of evaluation (we will not worry about precisely what elements comprise contexts and points of evaluation, except that contexts have as coordinates at least speakers, locations, times, and worlds, and points of evaluation have as coordinates at least worlds). Let us call functions from points of reference so understood to appropriate extensions *M-characters*. The *M-character* of a sentence, then, is a function from points of reference to truth-values. A sentence such as “I am here now” has an *M-character* that maps a point of reference  $\langle c, i \rangle$  to True iff the speaker of  $c$  is in the location of  $c$  at the time of  $c$  in the world of  $i$ . Thus, if we require, as is usual, that the speaker of  $c$  be at the location of  $c$  at the time of  $c$  in the world of  $c$ , the sentence will be true at all  $\langle c, i \rangle$  where the world of  $i$  is the world of  $c$ .<sup>17</sup> On the current conception of the semantics–pragmatics distinction, since *M-character* is the only unrelativized value assigned to expressions, the assignment of *M-character* to sentences exhausts semantics.

Our objection to this view is essentially that of Stalnaker (1999*b*). *M-characters* map context–index pairs to truth-values. But then *M-characters* don’t seem to be the kinds of things that are grasped in understanding sentences. Intuitively, it seems that in understanding a sentence, we combine the *referential contents* of the constituents of that sentence together according to its syntactic structure. Our understanding of a sentence in a context is due to a compositional procedure that calculates the content of the whole sentence from the referential contents of its parts. But on this conception of the semantics–pragmatics distinction, there is no room for a representation of this process, since sentences are not assigned referential contents in contexts at all. In short, this view does not assign *propositions* to sentences taken relative to contexts.

<sup>17</sup> This captures the idea that the sentence is “indexically valid”, while avoiding the result that the sentence expresses a necessary truth with respect to a given context  $c$ , since it will be false for many pairs  $\langle c, i \rangle$  as  $i$  is varied; we suppress the relevant definitions of *indexically valid* and *expresses a necessary truth with respect to context c*.

But propositions are needed to be the things grasped in understanding sentences in contexts. M-characters are not suited to play this role; nor are they appropriate objects of the attitudes or entities on which natural language operators operate. Thus, we think Stalnaker gave the right reason for rejecting this way of implementing this conception of the semantics–pragmatics distinction.<sup>18</sup>

On the other two ways of implementing the idea that semantics is distinguished from pragmatics by being concerned only with the assignment of non-relativized values to expression types, semantics consists of assigning to sentences entities that are, or determine, functions from contexts to contents or propositions (i.e., characters). On the first implementation, we assign to sentences (and complex expressions generally) *structured characters*. Assume that we assign to lexical items characters understood as functions from contexts to contents. The structured character of a complex expression is simply the concatenation of the characters of its lexical parts, where the characters are concatenated according to the syntactic structure of the complex expression. So for example, the structured character assigned to “I am hungry now” would be something like  $\langle I', H', n' \rangle$ , where  $I'$  is the character of “I”,  $H'$  is the character of “hungry” and  $n'$  is the character of “now”. This structured character, taken relative to the context  $c$ , yields the structured content or proposition  $\langle I'(c), H'(c), n'(c) \rangle$ .

The first thing to note is that on this conception of semantics, there is no non-trivial *semantic* composition. For on this view, the *semantic* content of a sentence is its structured character (sentences only have structured contents or express propositions relative to contexts; and so on the current way of understanding the distinction between semantics and pragmatics, these values that are relativized to contexts fall in the purview of pragmatics). But in deriving the structured character of a sentence like “I am hungry now”, one does not compose the characters in accord with the intuitively correct composition rules governing its syntactic structure. For example, one does not “saturate” the character of “am hungry now” with the character of “I” (or vice versa). As already mentioned, the character of “I am hungry now” is rather a result of concatenating the characters of the elements of the sentence. So, the semantic content of a sentence is not really determined by composing the semantic contents of the parts via non-trivial composition rules given by the sentence’s syntax. Instead, the semantic content of a sentence is determined by concatenating the semantic contents (i.e., characters) of the parts of the sentence. Eliminating non-trivial semantic composition, and in effect trivializing semantics, is an unattractive feature of this conception of the semantics–pragmatics distinction.

<sup>18</sup> The focus of Stalnaker’s criticism is that Montague’s system does not allow for the representation of propositions as values. However, Stalnaker did not appear to conceive of this as a criticism of Montague’s way of drawing the semantics–pragmatics distinction. In this sense, we are altering Stalnaker’s criticism to fit our target.

Our second concern about this way of understanding the semantics–pragmatics distinction is that we are skeptical about the utility of the sole semantics values, structured characters, allowed by such an approach. We see no use for these structured characters, and so do not see that a theory that employs them has any advantage over a theory that does not assign characters to complex expressions at all, but only assigns characters to syntactically simple expressions. If we are correct about this, then the current conception of the semantics–pragmatics distinction results in semantic values having no real use. Obviously, this would be a good reason to reject the current conception.

But are we correct in thinking that structured characters have no real use? David Braun (1994) has argued that we should accept a semantic theory that includes structured characters. Braun’s primary argument involves complex demonstratives. Braun considers the following two complex demonstratives:

- (1) that man
- (2) that man who is either talking to Bush or not talking to Bush

Braun assumes that the content of a complex demonstrative in a context is its referent. (1) and (2) have the same content in every context (assuming we hold all contextual factors, including associated demonstrations, speaker intentions, etc., constant—see Braun (1994, n. 7)). If we understand character functionally, as a function from context to content, (1) and (2) have the same character. But, Braun objects, (1) and (2) seem to differ in meaning. However, if we assign them only functional characters, we assign them no meanings on which they differ. Hence, Braun argues, we need to introduce structured characters, and assign different structured characters to (1) and (2). In this way, we can honor the intuition that (1) and (2) have different meanings.

Thus, Braun argues that, contrary to what we have claimed, there is an important use for structured characters. Before responding to Braun here, it is worth noting that he never actually gives an assignment of structured characters that assigns different structured characters to (1) and (2).<sup>19</sup> Given the incomplete nature of Braun’s proposal with respect to (1) and (2), it is somewhat difficult to evaluate.

Overlooking this shortcoming, we have two responses to Braun’s argument. First, even within the sort of framework Braun presupposes, on which the content or propositional contribution of a complex demonstrative in a context is its referent in that context (which we reject—see below), we are not convinced that honoring the pre-theoretical intuition that (1) and (2) have different meanings requires assigning them different structured characters. Even if the two phrases as a whole are assigned

<sup>19</sup> See e.g. the first paragraph of his section 11—the problem is that he never assigns structured characters to complex nominals like “man who is talking to Bush or not talking to Bush”.

no characters (characters only being assigned to their syntactically simple parts) or only functional characters (so that (1) and (2) get assigned the same character), surely the fact that one phrase contains words with contents (or characters) that the other phrase does not contain would be enough to explain normal speakers' intuitions that the phrases differ in meaning in some way. Indeed, Braun seems to admit this himself when he writes: "We judge that (1) and (2) differ in meaning at least partly because they have different meaningful parts."<sup>20</sup> So even within a framework of the rough sort endorsed by Braun, Braun has failed to give good reasons for assigning structured characters to (1) and (2).

Secondly, there are accounts of complex demonstratives that explain Braun's data without any appeal to character. For example, consider the quantificational account of complex demonstratives recently defended in King (2001). On this account, (1) and (2) are contextually sensitive quantifiers. As such, the predicative material in them gets contributed to propositions expressed in contexts by sentences in which they occur. Thus, (1) and (2) used in the same context make different contributions to propositions expressed in those contexts by sentences in which they occur. Obviously, the fact that (1) and (2) in any context have different contents (i.e., make different contributions to propositions) should explain why we have the intuition that they differ in meaning. So King's account of complex demonstratives straightforwardly accounts for the intuition that (1) and (2) differ in meaning without positing structured characters.<sup>21</sup>

<sup>20</sup> Braun (1994, 101). In a footnote to this remark, Braun says there are two other reasons for thinking that (1) and (2) differ in meaning. First, he says that to understand (2) one must grasp certain characters that one need not grasp to understand (1) (e.g. those of "talking", etc.). Secondly, he claims that (1) and (2) would receive different translations into other languages. But none of this requires positing structured characters for (1) and (2) either. One could claim that understanding (1) and (2) requires grasping the characters of their lexically simple parts and how those parts are syntactically combined. This would mean that understanding (1) requires grasping characters that need not be grasped in understanding (2), even if we have assigned no character to (1) and (2) as whole phrases (or only a functional character). Further, one could claim that it is a constraint on translation that one translate a syntactically complex phrase of one language into a syntactically complex phrase of the other languages that has the same syntax and parts with the same meaning as the phrase being translated (at least, when this is possible). Again, this explains why (1) and (2) would be translated differently and does not require assigning structured characters to (1) and (2). Finally, Braun also argues that intuitive differences in the meanings of certain dthat terms in an extension of English containing dthat terms requires assigning them different structured characters. However, the response we give here to the comparable claim about English complex demonstratives carries over straightforwardly to the claim about dthat terms in an extension of English.

<sup>21</sup> There are several other accounts of complex demonstratives that, like King's, account for the data discussed by Braun, without appealing to character. For example, according to Lepore and Ludwig (2000), complex demonstratives are similar to quantifier phrases in that the content of the nominal in a complex demonstrative affects the content of the whole sentence containing it. This theory, too, explains Braun's data without appeal to character. So too does the "appositive" account of complex demonstratives advanced by Joshua Dever (2001).

We conclude that Braun has given no compelling reason for positing structured characters. So we are left with the conclusion that on the way of distinguishing semantics from pragmatics we have been considering, where the sole job of semantics is to assign structured characters to complex expressions, the semantic values so assigned (structured characters) have no real function. Again, this is a significant problem for this approach.<sup>22</sup>

The final way to implement the view that semantics is distinguished from pragmatics in that the former is only concerned with assigning *unrelativized* semantic values to sentences (and complex expressions) is to take the job of semantics to be the assignment of *functional* characters to sentences (and complex expressions). On this way of proceeding, complex expressions are assigned as semantic values functions from contexts to appropriate contents. Further, the functional character assigned to a complex expression is compositionally determined from the functional characters of its simple parts and how they are syntactically combined.

To illustrate, consider the sentence “I am here”, and assume that as usual the characters of “I” and “here” are functions from contexts to the speakers of the contexts and the locations of the contexts, respectively. Loosely following Kaplan’s notation, regiment “I am here” as:

(3) Located (I, Here)

The character of “Located” will be a (constant) function from contexts to a two-place relation between individuals and locations (since we want structured contents, relations should be understood as relations-in-intension). Write “ $C(e)$ ” for the character of expression  $e$ . Then, suppressing considerable detail, the clause assigning a character to (3) (assuming the characters of “I” and “Here” mentioned above) loosely runs as follows:

(4)  $C(\text{“Located}(i, p)\text{”})$  for “ $i$ ” an individual term and “ $p$ ” a position term =  $f$  such that for any context  $c$ ,  $f(c) = \langle C(\text{“Located”})(c), \langle C(\text{“}i\text{”})(c), C(\text{“}p\text{”})(c) \rangle \rangle$

So the character of (3) is a function  $f$  that maps a context  $c$  to the structured proposition  $\langle L, \langle s, o \rangle \rangle$ , where  $L$  is the relation that the character of “Located” maps every context to, and  $s$  and  $o$  are the speaker and location of the context  $c$ , respectively. Note that we cannot complain here, as we did on the previous account, that there is no real compositional semantics. Here there is: the characters (semantic

<sup>22</sup> One argument we know of that structured characters have some important role to play was given in Richard (1983). Richard argued that structured characters (which he calls *meanings*) have a role to play in the semantics of belief-ascriptions. We don’t think such an account of the semantics of belief-ascriptions is promising, and Richard himself no longer endorses the position he defended in this paper (see Richard (1990, 122 n. 8)).

values) of parts of a sentence combine in accordance with the syntax of the sentence to give the characters (semantic values) of the whole.

So what is wrong with this view? Our first concern is that such an account intuitively assigns the wrong semantic significance to syntactic combination.<sup>23</sup> This fact is obscured somewhat by clause (4). For clause (4) leaves implicit the semantic significance that it accords to the syntactic concatenation of the elements of “Located(*i, p*)”. So let us consider how the view would work with a simpler example. On this view, the significance of syntactically combining, for example, a name and a simple predicate, say “Njeri runs”, is not that of predicating a property of an object. For the compositional semantics for this sentence must combine the *character* of “Njeri” with that of “runs” to yield the semantic value (character) of the whole sentence. Thus the semantic significance of the syntactic concatenation here is that of a function that maps a pair of a function *f* from contexts to objects and a function *g* from contexts to properties to a function *h* that maps a context *c* to the pair  $\langle f(c), g(c) \rangle$ . This seems grossly implausible as the meaning of the syntactic relation of predication.

Furthermore, it simply does not seem that in understanding such a sentence in a context, speakers employ this compositional semantics to determine a character for the sentence, and then apply it to the context yielding the relevant structured proposition. Intuitively, it seems rather that speakers evaluate in a context the characters of syntactically simple expressions in a sentence, and compositionally combine the resulting referential contents in grasping the proposition expressed by the sentence. Thus, the account seems to get the phenomenology of linguistic understanding wrong.

However, the most important objection to this account is that (phenomenological considerations of the sort just raised aside) there seems to be no empirical difference between this account and one that assigns characters only to syntactically simple expressions. After all, the job of character is to give us content, and we can assign contents to complex expressions in contexts using only the characters of the parts, and combining the contents they determine in those contexts. Thus, imagine a theory that assigns no character to the sentence (3) as a whole, but assigns it a content in every context as follows:

- (4') The content of “Located (*i, p*)” in any context *c* is  $\langle C(\text{“Located”})(c), \langle i^*, p^* \rangle \rangle$ , where *i*<sup>\*</sup> is the content of “*i*” in *c* and *p*<sup>\*</sup> is the content of “*p*” in *c*.

(We presuppose a recursive assignment of referential contents in contexts to complex individual and position terms—in the general case “*i*” and “*p*” could be syntactically complex in 4'.) As did (4), (4') assigns to the sentence (3) relative to a context *c* the

<sup>23</sup> The concern we develop here applies equally to the first way of implementing this conception of the semantics–pragmatics distinction.

content  $\langle L, \langle s, o \rangle \rangle$ , where  $L$  is the relation that the character of “Located” maps every context to, and  $s$  and  $o$  are the speaker and location of the context  $c$ , respectively.

The point is that both a semantics that assigns characters to simple expressions and recursively assigns characters to complex expressions *and* a semantics that assigns characters to only simple expressions allow for an assignment of the same contents in contexts to simple and complex expressions. So unless the functional characters of complex expressions have some *additional* job to do, they are unnecessary. But there seems to be no such additional job. Thus, on the present proposal, the only semantic value that is assigned via semantic composition, i.e. the characters of complex expressions, have no purpose. We could just as well do everything we need to without them. So here again, the present conception of semantics makes semantics out to be something that assigns only values with no real use to complex expressions. This, we think, is sufficient to dismiss the account.

In this section we have discussed three different ways of implementing the first conception of the semantics–pragmatics distinction, which stipulates that semantic values must be context-independent. We have found difficulties with all three ways of implementing this approach. This suggests the need for some semantic values to be relativized to contexts. But this insight is consistent both with our favored way of drawing the semantics–pragmatics distinction, and with the third way of drawing the distinction, according to which only a highly restricted range of context-sensitive expressions count as having genuine context-relative semantic values. In the next section we turn to a criticism of this third way of drawing the semantics–pragmatics distinction.

### III Semantic Content as Minimally Context-dependent

On the first way of drawing the distinction between semantics and pragmatics, semantic content is context-independent semantic value. On our favored way of drawing the distinction between semantics and pragmatics, there is a level of semantic value, which we are calling semantic content, that is relativized to context. One and the same expression can have different semantic contents in different contexts. On this conception of semantics, relative to a context, every indexical and demonstrative element in a sentence has potentially different *semantic* contents in different contexts. For example, the semantic contents of both “I” and “she”, relative to different contexts, may be different.

Some philosophers, however, hold that our favored conception of the semantics–pragmatics distinction allows context to affect semantic content to a greater degree than is plausible. We begin this section by discussing two worries that may motivate rejecting our favored conception in favor of the third, and more restrictive, conception of the semantics–pragmatics distinction discussed in the first section.



We then turn to an extended discussion and critique of this third conception of the semantics–pragmatics distinction.

The first worry one might have with our favored conception of the semantics–pragmatics distinction that might move one towards the third conception of the semantics–pragmatics distinction is that it does not allow for a distinction between speaker meaning and semantic content. For example, there is a distinction between the semantic content of a sentence on an occasion, and what the speaker meant by uttering that sentence on that occasion. Suppose John is confronted by a nice man drinking water, who John falsely believes to be drinking a martini. John utters the sentence “The man with the martini is nice”. The semantic content of John’s sentence on that occasion concerns the property of being a man drinking a martini, rather than the man to whom he intends to refer. However, John’s referential intentions concern the man he sees drinking water, so his speaker meaning is a proposition about that particular man. One might worry that by allowing speakers’ referential intentions to affect semantic content, as one is likely to do in adopting this second conception of the semantics–pragmatics distinction, one blurs this important distinction.<sup>24</sup> Restricting semantically relevant contextual effects to “automatic” indexicals, in Perry’s sense, obviates this concern.

We are not persuaded by this concern. If speaker intentions are relevant for determining the semantic content of some indexical expressions, then there must be a principled distinction between the kinds of speaker intentions that are relevant to the fixation of semantic content and those that are not. It is important to be as clear as possible, in particular cases of context-sensitive expressions, which speaker intentions are semantically relevant, that is, relevant for the fixation of semantic content. However, we see no basis for skepticism about the possibility of distinguishing, in particular cases, those intentions that are semantically relevant from those that are not.<sup>25</sup>

<sup>24</sup> We say “likely”, because of Lepore’s point, mentioned above, that one may adopt our preferred way of distinguishing semantics from pragmatics, but adopt a thoroughgoing anti-intentionalism about demonstratives and indexicals generally, of the variety favored by Chris Gauker. However, we think it very likely that speaker intentions do play a role in determining the reference of many indexical expressions.

<sup>25</sup> e.g., Jeffrey King, in his quantificational analysis of complex demonstratives, appeals to speaker intentions to fix the properties that saturate the argument places of the relation expressed by “that”. King is quite explicit (King (2001, 28–31)) about what sort of speaker intentions are relevant for the determination of content. Siegel (2002) supplies a particularly subtle discussion of the nature of those perceptual intentions that are relevant for reference fixing. Siegel’s discussion makes significant progress towards a criterion to distinguish between those sorts of intentions that are relevant for reference fixation, and those that are not. (We should note a particular debt to Maite Ezcurdia for discussion of these issues; her work bears directly on drawing the required distinctions between semantically relevant and semantically irrelevant speaker intentions.)

A second concern one may have with our favored conception of the semantics–pragmatics distinction is that allowing context to affect semantic content always brings with it a cost in complexity and systematicity to the semantic theory. This is a rather imprecise concern, but nevertheless, we suspect, a quite influential one. The concern is that there is an inverse correlation between the amount of semantic context-sensitivity a semantic theory recognizes and the systematicity and simplicity of that semantic theory. This concern would lead one to adopt the third conception of the semantics–pragmatics distinction over our favored conception (and taken to its limit, the first conception over the third).

We suspect this concern reflects a long-standing prejudice from earlier debates in the philosophy of language. For example, consider the stalking horses of Grice’s “Logic and Conversation” (Grice (1989a)). These philosophers (Grice’s “informalists” and “formalists”) held that the natural language counterparts of the logical expressions were so laden with context-dependency and vagueness in their use, that any “logic” of these natural language expressions would be “unsimplified, and so more or less unsystematic” (1989a, 24).<sup>26</sup> In contrast, Grice’s favored approach is to provide pragmatic explanations of context-dependency, whenever possible, presumably leaving the semantics “simple” and “systematic”.

We suspect that this dialectic has led philosophers of language, whether knowingly or not, to adopt the position that preserving the simplicity and systematicity of a semantic theory for a natural language somehow requires giving pragmatic (non-semantic) explanations of context-dependency whenever remotely possible. But to adopt this position is to accept the view that giving a context-dependent semantic treatment of a linguistic construction is tantamount to treating that construction as resistant to “systematic” semantic analysis. In the intervening decades between Grice’s work and today, there have been analyses of a host of constructions (such as modals, indicative conditionals, quantifier phrases, and adjectival modification, to name but a few) that incorporate context-dependence into the semantics without compromising the systematicity of the semantic theory (in a sense, this was also the purpose of Montague’s work on formal pragmatics). The view favored by Grice’s opponents concerning the ramifications of semantic context-dependency looks therefore to have been superseded by later developments.

In fact, we believe an even stronger response is possible to the view that semantic context-dependency is correlated with less simple and more unsystematic semantic theorizing. For it can turn out that claiming that certain sorts of context-dependency are not semantic can result in a less simple and more inelegant semantic theory. This point is the basis of our criticism of the third way of distinguishing semantics from

<sup>26</sup> Although in Grice’s writings these philosophers do not explicitly use the phrase “context-dependence” or “vagueness”, this is one plausible reading of what is meant by “unsystematic”.

pragmatics. By relegating certain contextual effects to pragmatics, this third way of forging a distinction between semantics and pragmatics results in a more complex and inelegant semantic theory.

Recall that on this third conception of the semantics–pragmatics distinction, among context-sensitive expressions only the content of automatic indexicals affects semantic content (relative to a context). So, automatic indexicals are the only semantically context-dependent expressions, on this conception. Non-automatic context-sensitive expressions do not contribute their contents to the semantic content (relative to a context) of sentences containing them. So, the semantic content of the sentence “I am human” is, relative to a context, the proposition, about the speaker of that context, that she is human. In contrast, the semantic content of the sentence “She is human” is not the proposition, about the demonstrated woman, that she is human. The reason that the semantic content of “She is human” is not the proposition about the demonstrated woman to the effect that she is human, on this third conception, is that the pronoun “she” is not an automatic indexical, and so its content in a context does not contribute to the semantic content of sentences containing it.

The first point to note is that this third conception of the semantics–pragmatics distinction is quite similar to the first conception. For the list of automatic indexicals is highly restricted, confined to words such as “I”, “today”, and “tomorrow”.<sup>27</sup> The vast majority of context-sensitive expressions are intentional. So semantic theories in accord with the first and third conceptions of the semantics–pragmatics distinction share considerable overlap in their assignment of semantic content. As a result, the third conception of the semantics–pragmatics distinction inherits many of the concerns facing the first. But there are additional problems with it besides.

On this third conception of the semantics–pragmatics distinction, the semantic content of “I am human”, relative to a context, is the proposition, concerning the speaker in the context, that she is human. But the semantic content of “She is human”, relative to a context, is not a proposition. The English pronoun “she” perhaps does possess some context-independent meaning; the referential content of “she” relative to a context must be female. So it seems one plausible semantic content for “She is human”, relative to a context, is the property of being a female human, a function from a possible situation to the set of women in that situation (or a “propositional radical” that determines such a function).

But this is, *prima facie*, a worrisome result. The sentences “I am human” and “She is human” differ only in that the first contains the first-person pronoun in subject

<sup>27</sup> As Robyn Carston (ms) has pointed out, even indexicals such as “now” and “here” are not really automatic, since speaker intentions determine the temporal scope of the referential content of “now” relative to a context, and the spatial dimensions of “here” relative to a context.

position, and the second contains the third-person pronoun (and the concomitant facts about case). As far as their structural features are concerned, the sentences are identical. So the mapping rules from sentences to semantic contents should be the same. But the sentence “I am human”, relative to a context, is supposed to express a proposition, a function from possible situations to truth-values, while the sentence “She is human”, relative to a context, expresses a property, namely, the property of being a female human, a function of a very different sort. On this conception of the semantics–pragmatics distinction, if the subject term of a sentence is non-indexical, or an automatic indexical, then, roughly speaking, it contributes an object to the semantic content of that sentence relative to a context, but if it is a non-automatic indexical or demonstrative, then it contributes a property that is conjoined with the property expressed by the predicate of that sentence. It is not immediately clear how to derive these results in a non-*ad hoc* fashion.

We have focused, for simplicity’s sake, on very simple sentences in making this point. But it should be clear that the point generalizes to a host of different constructions. For example, according to this third conception of the semantics–pragmatics distinction, (5) and (6), despite sharing the same syntactic structure, have semantic contents of drastically different types:

- (5) Every woman who met me yesterday is smart.
- (6) Every woman who met her yesterday is smart.

(5), relative to a context, expresses a proposition. (6) does not express a proposition, even relative to a context, because it contains the non-automatic indexical “her”. Even more problems are involved in constructing the semantic value of a sentence like “That man in her car is wearing a hat”, where the complex demonstrative “that man in her car” involves a kind of double incompleteness. What kind of semantic values such sentences express, according to this third conception of the semantics–pragmatics distinction, and how to construct non-*ad hoc* semantic rules that assign these semantic contents to these sentences, is an interesting technical question.<sup>28</sup>

It is no doubt possible to write some kind of semantic algorithm to do the trick of assigning distinct types of semantic values to (5) and (6), relative to a context. However, making semantic content sensitive to the effects of context on

<sup>28</sup> Thomas Hofweber has made the following suggestion to us. One could add, to each domain of type T, a set of “incomplete” entities of that type. So, in addition to complete properties, there would be incomplete properties, and in addition to complete propositions, there would be incomplete propositions. Composition rules could, for example, combine an incomplete object with a complete transitive verb meaning to yield an incomplete property. This move seems *ad hoc* to us. One wonders with what right one can classify an “incomplete X” as an X at all. Furthermore, one wonders how one would marry this semantics with the notion of “pragmatic enrichment”. How would interpreters “add” on constituents to meanings, thus construed?

the content of all context-sensitive lexical items, as on the second conception of the semantics–pragmatics distinction, results in simpler semantic theory. We can draw a moral here, one that tells against the view with which we began concerning the allegedly disruptive effects of contextual sensitivity on the systematicity of semantics. It can indeed turn out that allowing certain contextual effects to have semantic import allows one to give a simpler semantics than would be otherwise possible. Whether there is in the end a compelling “simpler semantics” argument against advocates of this third conception of the semantics–pragmatics distinction depends, of course, upon the details of an as yet non-existent semantic theory that is in accord with it.

There is another set of worries that faces the advocate of this third conception of the semantics–pragmatics distinction who adopts the structured proposition conception of semantics. On this method of semantic interpretation, semantics takes place in two stages. First, there is an algorithm assigning structured propositions to sentences relative to a context. The resulting structured entity is a concatenation of the semantic contents of the elements of the sentence. The second stage of semantic interpretation involves giving a truth-definition for the resulting structured propositions. At this stage in the semantic process, the elements of the structured proposition are combined in accord with semantic composition rules.<sup>29</sup>

<sup>29</sup> Things are actually a bit more complicated than we make them out to be here. We talk here as though *no* significant composition of semantic contents of the elements of a sentence occurs in the mapping from sentence to proposition. This is true on the account of structured propositions sketched in King (1995, 6), on which the structures of propositions are identical to the structures of sentences expressing them (in a context). Thus, the semantic contents (in the context) of the lexical items in the sentence are not composed in any significant way in the mapping to the proposition: the proposition is literally the concatenation of these semantic contents, and is structurally identical to the sentence expressing it. By contrast, in the mapping from propositions to truth-values, there is significant composing of the constituents of the proposition (i.e. the semantic contents of the lexical items in the sentence expressing the proposition) to yield a truth-value. However, one might hold that in the mapping from sentence to proposition for example, a complex predicate contributes a property or a property and a time to the proposition (this view is suggested by clauses (22) and (28) of Salmon (1986, 145–6)). On this view, the semantic contents of the lexical items in the predicate apparently *are* semantically composed to yield the property contributed by the predicate to the proposition (this would require a recursive assignment of properties to complex predicates, which Salmon does not actually provide—see again clause (28)). On this view, then, some significant semantic composition occurs in the mapping from sentence to proposition and the rest occurs in the mapping from proposition to truth-value. How much composition occurs in the first-stage mapping from sentence to proposition depends on how much of the sentence structure is “preserved” in the structure of the proposition. King (1995) argues that an account on which *no* significant composition occurs in the first-stage mapping is preferable. In any case, the more sentence structure is preserved in proposition structure, the less composition occurs in mapping from sentence to proposition, and the more occurs in the mapping from proposition to truth-value. The point remains that on any theory of structured propositions that preserves most of the structure of a sentence in the structure of the proposition, most of the significant composition of semantic contents occurs in the mapping from propositions to truth-values.

The advocate of this third conception of the semantics–pragmatics distinction might find this framework attractive. For example, she could argue that non-automatic indexicals contribute, say, an ordered pair of a property and an object representing a gap, such as the empty set, to the structured proposition expressed by sentences containing them.<sup>30</sup> The structured proposition assigned by the algorithm to a sentence is not a function of the composition rules corresponding to the syntactic structure of that sentence; the semantic values of the words in the sentence are simply concatenated in the structured meaning. So in the mapping to the structured proposition, one does not have to face the question of how to compose the semantic content of a non-automatic indexical with other elements. One could then argue that pragmatics intrudes before the truth-definition for structured propositions. In particular, pragmatics would intrude to assign values to the empty slots in the structured proposition.

On this picture, the semantic content of a sentence such as (6) would be a concatenation of properties and functions of various sorts, corresponding to the semantic values of the lexical items in (6). However, in the semantic content, corresponding to the place of the non-automatic indexical “her” in (6), would be an object representing a gap and the property of being a woman. Pragmatics would then somehow replace this ordered pair by a contextually salient woman. Then, together with the semantic composition rules corresponding to the syntactic structure of the sentence, one could calculate the truth-conditions of the resulting proposition.

However, we believe that the picture that results amounts, if anything, to a rejection of a genuine distinction between semantics and pragmatics, rather than a novel way to draw one. On the structured propositions picture defended by King (1995, 6), Soames (1987), and others, semantic interpretation is a two-stage process. First, there is the assignment of a structured entity to a sentence relative to a context. Secondly, there is the definition of truth for that entity. On the picture we are now considering, pragmatics intrudes between the first stage of this process and the second. That is, pragmatic processes enter in to “fill in” the empty slots in the structured entity.<sup>31</sup>

We believe that no one who thinks that there is an interesting and important distinction between semantics and pragmatics should be content with the resulting picture of semantics and pragmatics. For the view amounts to a rejection of the possibility of a strictly compositional semantics for a language containing non-automatic indexicals. In mapping the sentence to the structured entity, one does not appeal to the composition rules corresponding to the syntactic structure within the sentence;

<sup>30</sup> Braun (1993), following a suggestion by Kaplan, advocates a gappy structured meanings account for the semantic content of sentences containing names without bearers.

<sup>31</sup> Soames (2005) develops a picture like this, albeit not one clearly motivated by the distinction between automatic and non-automatic indexicals.

one simply concatenates the semantic values of the individual terms in the sentence. The composition rules expressed by the syntactic structure of the sentence are relevant only for the definition of truth for the resulting structured entities. But pragmatic mechanisms intervene before the truth-definition, on this picture. That is, the gaps in the semantic contents are “filled” by enrichment before the definition of truth.

So, on this picture, composition rules do not just compose semantic values; they compose semantic and pragmatic values. Indeed, on this view, in interpreting most sentences, composition rules do not compose only semantic values. Significant composition of the elements of structured propositions is then, strictly speaking, part of pragmatics, not semantics. Perhaps one way of describing this is as a rejection of the possibility of a strictly compositional semantics for a natural language. There is only an attenuated sense in which one is preserving a significant role for semantic content on such a way of dividing the labor between semantics and pragmatics.

According to the first conception of the semantics–pragmatics distinction, the semantic content of an expression was its context-independent semantic value, and every other kind of content associated with the sentence was pragmatic in nature. On our favored conception of the semantics–pragmatics distinction, there are two levels of genuinely semantic content. On the first level are properties of individual words, rather than complex expressions. This is standing meaning, or character. On the second is semantic content, a property of expressions relative to contexts. The semantic content of any term relative to a context is the referential content in that context that is in accord with its standing meaning, and the semantic content of a complex expression is derived by combining the semantic contents of its parts in accord with the composition rules corresponding to its syntactic structure. On the third conception of the semantics–pragmatics distinction we have been discussing in this section, only some context-sensitive terms contributed their contents in a context to the semantic content, relative to that context, of more complex expressions containing them.

We know of no compelling arguments against our favored way of drawing the semantics–pragmatics distinction. In contrast, we have raised concerns with the two natural competing conceptions. Some of these concerns may be met by greater attention to the technical details of a semantic theory that respects, say, the third conception of the semantics–pragmatics distinction. But until such a theory is advanced, in the absence of any compelling objections against our own favored conception of the semantics–pragmatics distinction, it should clearly be the default conception of the relation between semantics and pragmatics.

On this conception, the effects of context on semantic interpretation are restricted to what we have called above “weak pragmatic effects”. That is, context can only affect the semantic interpretation of an expression by being involved in the interpretation of some constituent of that expression. But semantic content is not limited to denotation

assignment to “automatic” indexicals. The content of any syntactic constituent is relevant to the semantic content of expressions containing it. This is the picture of semantic content we will assume in the remainder of this chapter.

#### IV Skepticism and Modesty

Some disagreements in the philosophy of language concern the correct conception of the semantics–pragmatics distinction. No doubt some of these disagreements are terminological. But perhaps some are not. For example, in the last two sections we have suggested that there are some ways of drawing the distinction between semantics and pragmatics that face objections that our preferred conception does not. This suggests that there are some non-terminological issues involved in deciding what the “correct” distinction between semantics and pragmatics is. Be that as it may, once one fixes upon one way of drawing the distinction between semantics and pragmatics, many obviously non-terminological disputes arise.

We have settled upon one favored way to draw the semantics–pragmatics distinction. In this section we use this semantics–pragmatics distinction to describe three different views about the relative importance of the semantic content of a sentence, as we have characterized it. The disputes between these positions are definitely not terminological. In fact, they are largely empirical disputes, the resolution of which will require decades of cross-linguistic syntactic and semantic research, together with psychological studies. The different positions we describe correspond essentially to predictions about the ultimate outcome of these investigations.

The first position, and one which will centrally occupy us in the final section of this chapter, is what we may call *semantic skepticism* about what is said. Semantic skeptics hold that at no stage in linguistic interpretation does the semantic content of a sentence play a privileged role. For example, what is accessible in linguistic interpretation is a level of content that is thoroughly affected by strong pragmatic effects, in the sense discussed in the first section. Indeed, according to the semantic skeptic, there is no reason for an account of language understanding to grant the semantic content of a sentence, in the sense we have adopted, any important role. Every interesting level of content is affected by strong pragmatic effects.

According to the semantic skeptic, individual words have semantic content. However, there is no interesting notion of semantic content for sentences. The attitude of the semantic skeptic towards the semantic content of a sentence is therefore much like the attitude we adopted in the second section towards the character of a sentence. It is certainly an entity that one can define. But it lacks an interesting or central role in an account of any significant natural language



phenomenon. Semantic skepticism has been one of the major positions among cognitive scientists for the last twenty years, thanks in large part to the influence of Relevance Theory, the theory of linguistic interpretation introduced and defended by Dan Sperber and Deirdre Wilson (1986). Its other prominent advocates include Robyn Carston, Stephen Levinson, and François Recanati, all of whom have been defending for years, on various grounds, the thesis that there is no interesting semantic notion of what is said by a sentence.<sup>32</sup>

The second position is what we call *semantic modesty* about what is said. In contrast to semantic skeptics, semantically modest theorists do believe that the semantic content of a sentence has a significant role to play in an account of linguistic understanding. Such theorists are loath to give up on a theoretically important notion of the semantic content of a sentence. However, they believe this role to be minimal. In particular, the semantically modest theorist agrees with the semantic skeptic that what is accessible to interpreters as the content of a speech act is thoroughly affected by strong pragmatic effects. Indeed, semantically modest theorists have produced some of the strongest arguments for this conclusion. However, the semantically modest theorist insists that the semantic content of a sentence still has a privileged role to play in an account of linguistic understanding. That is, the semantically modest theorist holds that there is an important notion of sentence content that is not sensitive to strong pragmatic effects.

Semantic modesty is a position that comes in degrees. On the one extreme, there are those who hold that most of our intuitions about what is said by a sentence in a context are affected by strong pragmatic effects. On the other end of the spectrum are theorists who hold that while generally intuitions about what is said by a sentence are reliable guides to semantic content, there is a restricted range of cases in which strong pragmatic effects may affect speaker intuitions about what they take to be the semantic content of a sentence. In what follows, we restrict the thesis of semantic modesty to the first extreme of this spectrum, since this is the position that has gained particular popularity among philosophers of language.

Three philosophers who have directly argued for an extreme version of semantic modesty are Kent Bach (1994, 2002a, b) and Herman Cappelen and Ernie Lepore (1997,

<sup>32</sup> It is true that Sperber and Wilson (1986, ch. 4) do talk of “sentence meanings”. However, for them, sentence meanings “do not encode thoughts”; they are rather “sets of semantic representations, as many semantic representations as there are ways in which the sentence is ambiguous” (193). Semantic representations, in turn, are “incomplete logical forms” (ibid.). So, sentence “meanings” are in fact not semantic content at all—they are sets of expressions. What plays a theoretical role in Relevance Theory is not any notion of semantic content for a sentence, but rather the *syntax* of a sentence—its syntactic logical form. It is this that is subject to pragmatic enrichment, to obtain the representation of the thought. So, on Sperber and Wilson’s theory, there is no role played at all by any notion of the *content* of a sentence, as opposed merely to its syntax.

2002).<sup>33</sup> A different group of philosophers of language are led to embrace semantic modesty indirectly, by some of their antecedent semantic commitments concerning certain linguistic constructions. Most salient in this regard are the *neo-Russellians*. Neo-Russellians have a certain view about the semantic content of propositional attitude sentences. The neo-Russellian position is that the semantic content of a sentence such as “John believes that Mark Twain was an author”, relative to a context, is that John bears the belief relation to the singular proposition, concerning Mark Twain, that he was an author. According to the neo-Russellian, this is the very same semantic content, relative to the same context, as the sentence “John believes that Samuel Clemens was an author”. However, neo-Russellians recognize that speakers intuitively believe that these sentences have different semantic contents, since speakers intuitively believe that their semantic contents can diverge in truth-value. So, neo-Russellians are committed to the view that there is a class of sentences with a semantic content that is relatively opaque to speaker intuitions. Neo-Russellians therefore find semantic modesty an immensely plausible position.<sup>34</sup>

The main worry with semantic modesty is its tendency to collapse into semantic skepticism. Advocates of semantic modesty expend their greatest efforts arguing for an error theory about ordinary speaker intuitions about semantic content. That is, just like advocates of semantic skepticism, semantically modest theorists are most eager to establish that what ordinary speakers grasp in a speech act is not the semantic content of the sentence uttered relative to that context, but is instead thoroughly infected by strong pragmatic effects. But rarely do semantically modest theorists bother to explain what privileged role they believe semantic content in fact plays in language understanding. So it ends up being somewhat of a mystery what role these theorists believe semantic content has in an account of language understanding.

For example, one reason theorists have in producing semantic theories is to explain the systematicity and productivity of language understanding.<sup>35</sup> Given a finite vocabulary, and grasp of the composition rules expressed by syntactic structures, speakers have the ability to grasp the propositions expressed by an infinite number of sentences. If language users employ a compositional semantic theory in grasping the

<sup>33</sup> One difference between the version of semantic modesty advocated by Bach and the version advocated by Cappelen and Lepore is that (as we have discussed above) Bach thinks that the semantic content of a sentence relative to a context can be a non-propositional entity, whereas Cappelen and Lepore retain the thesis that it must be a proposition, or (in their Davidsonian framework) a truth-condition.

<sup>34</sup> Of course, one can imagine a neo-Russellian who rejects full-blown semantic modesty, and restricts her views about the inaccuracy of speaker intuitions about semantic content to propositional attitude ascriptions. Such a neo-Russellian is, we take it, a target of Saul (1997). In recent years neo-Russellians have exhibited an attraction to the more extreme forms of semantic modesty.

<sup>35</sup> Thanks to Ernie Lepore for discussion that led to this point.

contents of speech acts, then one has a satisfactory explanation of the systematicity and productivity of a speaker's grasp of an infinite number of novel utterances. For then one can explain a language user's grasp of what is expressed by the utterance of a novel sentence by appealing to the fact that she grasps the words in the sentence and their modes of combination, together with whatever contextual information is required to interpret the context-sensitive elements in the sentence. Given compositionality, nothing else is required to explain her grasp of the proposition expressed, since what is expressed by the utterance is then a function of what she already grasps.

But it is difficult to see how the semantically modest theorist can appeal to this justification for taking the project of semantic theory seriously. For the semantically modest theorist holds that, in ever so many cases of successful communication, when someone utters a sentence, her interlocutor grasps a proposition that is not determined compositionally from the values of the elements in the sentence and their composition. Rather, tacit unsystematic pragmatic processes intervene. This is, after all, why competent hearers are so often wrong about semantic content, according to the semantically modest theorist. So the semantically modest theorist is committed to an alternative explanation of our grasp of an infinite number of novel utterances, one that does not proceed by attributing our competence to a simple, compositional mechanism. So it is hard to see how she could accept arguments for the necessity of semantic theory that proceed on the assumption that grasp of a compositional semantic theory is the only way to explain our ability to grasp an infinite number of novel utterances given finite means.

According to the semantically modest theorist, there is some role to be played by their minimal notion of the semantic content of a sentence, relative to a context. This role cannot be played by the semantic contents of words relative to a context, together with their syntactic structures. The difficulty of distinguishing between semantic skepticism and semantic modesty is that it is hard to see what explanatory role could be played by such a minimal notion of the semantic content of a sentence relative to a context, that could not also be played solely by appeal to the semantic contents of words (relative to contexts), together with the syntactic structure of sentences. But unless there is a genuine explanatory role that is played by the semantically modest theorist's notion of sentence content, relative to a context, semantic modesty dissolves into semantic skepticism.

A final worry about semantically modest theorists is that their views threaten to undermine the very data for semantic theories. As Stanley and Szabó note, "accounting for our ordinary judgments about the truth-conditions of various sentences is the central aim of semantics. Since these judgments are the data of semantic theorizing, we should be careful with proposals that suggest a radical revision of these judgments" (p. 90). So, semantic modesty not only obscures the

purpose of the semantic project, but also removes the central empirical data for its claims. For if, as semantic modesty has it, speaker judgments about truth-conditions are not reliable guides to the semantic content of sentences, it becomes unclear how to evaluate semantic proposals (cf. also DeRose (1995, section 16; 1999, sections 9 and 10)).

So, from our perspective, both semantic modesty and semantic skepticism pose similar threats. First, it is difficult to see what the purpose of semantic theory should be, on either view. For according to these views, what is intuitively said by an utterance of a sentence is not significantly constrained by the syntactic structure of that sentence, since the result of semantically interpreting the syntactic structure of a sentence relative to a context far underdetermines what is intuitively said by that utterance. Secondly, one may worry about the explanatory value of accounts of what is intuitively said by an utterance that appeal to pragmatic processes unconstrained by syntax. Given the degree to which pragmatic mechanisms can affect what is intuitively said on these views, the sense in which such theories make explanatory predictions about intuitive judgments is unclear. For example, one worry is that the unconstrained nature of such accounts leads to an unconstrained over-generation of predictions about what a sentence can be used to intuitively say, relative to a context (see Chapter 5). In contrast, an account of what is intuitively said by a sentence relative to a context of the sort advocated here and in Chapter 1, that appeals only to syntax, semantics, and weak pragmatic effects, does not involve appeal to mysterious unconstrained processes. So there needs to be a rather strong justification for the view of semantic skeptics and semantically modest theorists that what is intuitively said is not so constrained.

The kind of arguments such theorists advance all take the following form. First, they produce a linguistic construction *L*, uses of which speakers intuitively believe to have a certain interpretation *I* relative to a context *c*. Then, they argue that the interpretation *I* cannot plausibly be the semantic content of *L* in *c*, because *I* is not just the result of assignment of referential contents in *c* to the parts of *L* and composition. That is, proponents of semantic skepticism and modesty argue that the relevant interpretation is the result of strong pragmatic effects. If so, then speaker intuitions about semantic content are sensitive to strong pragmatic effects.

Opponents of such arguments have in recent years argued that the examples discussed by semantically skeptical and semantically modest theorists are insufficient to support their conclusions.<sup>36</sup> For example, Stanley and Szabó (Chapter 2) argue, as against advocates of semantic skepticism and semantic modesty, that our intuitions about what is said by quantified sentences track the semantic content of these sentences, and additional work extends these points to a number of different constructions that

<sup>36</sup> See Chapters 1, 2, 3, 5, and Szabó (2001).

semantically skeptical and semantically modest theorists have taken to show that speaker intuitions about semantic content are not reliable.<sup>37</sup>

In the next section we challenge one powerful argument in favor of semantic skepticism or modesty, one that has not yet received a response. Versions of the argument have been around for several decades. Perhaps the simplest form of the argument involves negated sentences such as:

(7a) John doesn't have three children, he has four.

(7b) I'm not happy, I'm ecstatic.

Assuming, as is usual, that "three" means at least three, in these cases, the negative element appears to negate not the semantic content of the sentence, but rather the semantic content together with its implicature. For example, a pragmatic (scalar) implicature of "John has three children" is that John has no more than three children. Intuitively, what is negated in (7a) is the proposition that John has three children and no more than three. If so, then such cases are examples of strong pragmatic effects on the intuitive content of the speech act, and therefore support an error theory of speaker intuitions about semantic content, in the sense of "semantic content" we have adopted. Let us call this alleged kind of strong pragmatic effect, a "pragmatic intrusion".

There are problems with the pragmatic intrusion account of the examples in (7). Larry Horn has convincingly argued that the examples in (7) are instances of "metalinguistic negation", where one is denying the appropriateness of an utterance, rather than its content, and even provides compelling phonetic diagnostics for such uses of "not". Horn's discussion undermines any clear argument from examples such as (7) to an error theory of speaker intuitions about semantic content. For example, it shows that it is wrong to characterize these cases as ones in which what is negated is semantic content somehow augmented with pragmatic implicatures. So defenders of pragmatic intrusion have wisely chosen not to stake their case on these uses of "not".

<sup>37</sup> The debate surrounding attitude-ascriptions is more complicated. One basic problem is that different parties in the debate do not share the same conception of semantic content. For example, Crimmins and Perry (1989) and Crimmins (1992) argue, apparently in response to the neo-Russellian, that modes of presentation affect the semantic content of propositional attitude ascriptions. But Crimmins and Perry have a different conception of semantic content from the average neo-Russellian. According to Crimmins and Perry, there are "unarticulated constituents" of the semantic content of some sentences (an unarticulated constituent of an utterance *u*, in the sense of Crimmins and Perry, is an element of the semantic content of *u* that is not the value of any element in the sentence uttered). Like us, most neo-Russellians accept a conception of semantic content according to which, by definition, there are no unarticulated constituents of semantic contents. So the substance of the debate between Crimmins and Perry, on the one hand, and neo-Russellians, on the other, concerns the right definition of semantics, rather than attitude constructions. Peter Ludlow (1995) has, however, argued quite directly for the syntactic representation of modes of presentation, which meets the neo-Russellian more squarely on her own ground.

In recent years, however, increasingly more sophisticated arguments for pragmatic intrusion have made their way into the literature. For example, Robyn Carston (1988) has, with the use of these more sophisticated arguments, tried to provide evidence for the existence of the non-Gricean, relevance-theoretic pragmatic process of *explicature*. More recently Stephen Levinson has argued that the sorts of examples discussed by Carston support a related notion, what he calls *implicature intrusion*. Both explicature and implicature intrusion, if genuine, would entail the existence of strong pragmatic effects on the intuitive content of speech acts.

The debate between Carston and Levinson is important for pragmatics, but not for our purposes. Our expression “pragmatic intrusion” is intended to be neutral between notions such as that of the relevance-theoretic explicature, Recanati’s free enrichment, and the more Gricean notion of implicature intrusion. If there is a significant range of cases of pragmatic intrusion, then that would provide strong support for the existence of strong pragmatic effects on what is intuitively said. In the next section, focusing on Carston and Levinson’s arguments for pragmatic intrusion, we argue that the case for pragmatic intrusion is not compelling.

## V Pragmatic Intrusion

Levinson himself is admirably cautious about his arguments for pragmatic intrusions. As he notes, “There will always be doubts about whether a better semantic analysis of the relevant construction might not accommodate the apparent pragmatic intrusions in some other way” (Levinson (2000, 214)). In this section we intend to substantiate Levinson’s concern, by showing that the existence of pragmatic intrusion has yet to be substantiated. The arguments for it in the literature rest upon an inadequate grasp of the syntax and semantics of the particular constructions that appear to give rise to it.

The first case of pragmatic intrusion discussed by Levinson involves the comparative “better than”. The first set of examples, discovered by Wilson (1975, 151), involve sentences like the following:

- (8) Driving home and drinking three beers is better than drinking three beers and driving home.

As Wilson notes, it would seem that what is said by (8) is that it is better first to drive home, and then to drink three beers, than first to drink three beers, and then to drive home. If Grice is correct that “and” does not have a temporal meaning, then the temporal information has to come from a pragmatic effect. But then what is said by (8) appears to be sensitive to the pragmatics of the use of the clauses flanking

“better than”. If so, then there are strong pragmatic effects on the intuitive contents of speech acts.

But as the discussion in and of Cohen (1971) makes clear, appealing to the data about the temporal nature of conjunction is complex. According to most twenty-first-century syntactic theories, each clause contains a tense phrase, the contents of which give information about the time that the proposition expressed by the clause concerns. So in a sentence such as “John came in and sat down”, there is a tense phrase associated with each verb in the underlying syntactic structure. Relative to a context, this conjunction expresses a proposition of the form “John came in at time  $t$  and sat down at  $t + n$ ”. As Levinson sums up the situation, for all the data shows, “Grice’s position with respect to ‘and’ could be maintained, while the truth-conditional nature of the ‘and then’ inference could be attributed to the implicit indexical sequence of reference times in the two verbs” (Levinson (2000, 199)).<sup>38</sup>

However, the conjuncts in (8), as Levinson notes, do not involve finite clauses, and hence have no obvious tense. Furthermore, (8) is a generic sentence, which adds an additional factor to the interpretation of the data. However, we think that the same response is available, despite these additional complexities. For many non-finite clauses do carry clear temporal information. For example, consider:

- (9a) John will remember bringing the cake.
- (9b) John remembered to bring the cake.

(9a) can express the proposition that John will remember bringing the cake *today*. (9b) can express the proposition that John remembered to bring the cake *today*. Such evidence suggests the existence of an indexical temporal element in the syntax of some gerunds and infinitives, which, relative to a context, can be assigned a contextually salient time or event.<sup>39</sup> Furthermore, the infinitives and gerunds in (9) can be modified by explicit temporal adverbs such as “today” and “tomorrow”. On virtually all contemporary accounts of the functioning of such adverbs, they modify the values of syntactically represented times or events.<sup>40</sup>

<sup>38</sup> The fact that the temporal element in the second conjunct is later than the temporal element in the first conjunct is due to a pragmatic maxim governing the description of events; as Partee (1984, 254) describes the condition, “there is a past reference time  $r$ - $p$  specified at the start of the discourse, and . . . the introduction of new event sentences moves the reference time forward”. We will argue that this is a pragmatic maxim that affects semantic content by influencing the semantic content of temporal elements in the syntax.

<sup>39</sup> We are in fact agnostic here as to the precise nature of the contextually sensitive temporal element here (though see King forthcoming). For example, in Higginbotham’s (2002) account of sequence of tense, tenses express relations between events. On Higginbotham’s account, times would not be the values of the syntactically represented context-sensitive elements.

<sup>40</sup> See Enc (1987, 652); King (forthcoming); Ogihara (1996, 41–9, 56–60). Parsons (1990, ch. 11) is another obvious example, though whether Parsons accepts the letter of this point depends upon

There is also additional clear syntactic evidence in favor of the existence of a temporal element in many infinitives and gerunds. For example, this element can interact with higher operators in variable-binding configurations, as in:

(10a) John asked to bring a cake many times.

(10b) John tried singing many times.

The sentences in (10) are ambiguous. On one reading of (10a), John asked something many times; on this reading, the temporal quantifier “many times” binds a temporal element associated with the verb “ask”. But on the second reading, John asked once, at some contextually salient past time, to bring a cake many times. On this reading, the temporal quantifier “many times” binds a temporal element associated with the infinitive “to bring a cake”.<sup>41</sup> Similarly, there is a reading of (10b) according to which John tried once in the past to sing many times. Such evidence provides support for the existence of a temporal element associated with many infinitives and gerunds (see Chapters 1 and 3 for discussion of this strategy).<sup>42</sup>

There is thus good evidence for a temporal element in the syntax of many non-finite clauses. (8) is a generic sentence, and so involves an unselective generic quantifier.<sup>43</sup> The temporal elements in the non-finite clauses in (8) are bound by this generic operator, yielding the interpretation in (11):

(11) Gen- $x, y$  [Better-than( $x$  driving home at  $y$  and  $x$  drinking three beers at  $y + 1$ ),  
( $x$  drinking three beers at  $y$  and  $x$  driving home at  $y + 1$ )].

(11) is the semantic content predicted by a standard theory of generics, coupled with a plausible hypothesis about the syntax of non-finite clauses. Thus, the right syntax and semantics for constructions such as (8) simply predicts they have the readings

the relation between his “subatomic” representations and syntax proper. Discourse Representation-Theoretic accounts of temporal adverbs also treat them as setting descriptive conditions on syntactically represented temporal elements, though in the case of such theories, they are represented syntactically only in the level of syntax called Discourse Representation Structure (see Partee (1984)).

<sup>41</sup> It does not matter for our purposes whether this temporal element is in argument or adjunct position; we do not here commit ourselves on this matter. If the temporal element is an adjunct, perhaps its presence is optional.

<sup>42</sup> There is now a rather large and intricate literature in syntax on the question of tense in infinitives and gerunds, starting with Stowell (1982). Infinitives and gerunds that must inherit their tense properties from the matrix clause, and hence do not allow temporal modification, are usually not thought to contain a tense phrase (or to contain one that is deficient in its case-assigning properties). The gerunds in examples such as (8) are not contained in embedded clauses. However, they clearly can be modified by temporal adverbs such as “today” and “at nine o’clock”. Again, this alone suggests that they contain a temporal element.

<sup>43</sup> For unselective binders, see Lewis (1998) and Heim (1982, ch. 2). An unselective binder binds multiple variables in its scope.



at issue. No appeal to effects of implicature on intuitive content are required in this explanation.

Suppose a theorist is in the position of arguing that the semantic content of a certain construction must be augmented by pragmatic information to account for natural readings of that construction. To make out this claim, it is incumbent upon the theorist to provide a sketch of the correct syntax and semantics of the relevant construction. For it is only then that one can evaluate the claim that the syntax and semantics for the relevant construction does not by itself deliver the reading at issue. Far too often, those who advocate the thesis that a certain reading of a construction is due to pragmatics rather than semantics fail to live up to their obligation, and so make claims that are highly speculative and hence difficult to evaluate. Before we can evaluate the other alleged examples of pragmatic intrusion, we need to sketch a semantics for “better than”, with respect to which we can evaluate the claim that one cannot come up with a syntax and semantics that deliver the right readings.

For the sake of simplicity, we will treat only the “better than” relation that relates propositions.<sup>44</sup> One natural proposal for the semantics for “better than”, as a relation between propositions, is as in (12):

- (12) Better-than ( $p, q$ ) if and only if the most similar  $p$ -worlds are preferable, in the contextually salient sense, to the most similar  $q$ -worlds.

There are two dimensions of contextual sensitivity required for this semantics. First, the relation of similarity between possible worlds that is relevant for the truth-conditions of a “better than” sentence in a context is determined in part by facts about that discourse context. Secondly, the sense of preferability at issue is also a function of the discourse context.<sup>45</sup>

We do not mean to advocate this as a semantics for “better than”; we find it unpleasantly skeletal. But, as we shall show, one can account for the other alleged examples of pragmatic intrusion without abandoning it. That is, as far as we have seen, the semantics just described can account for all of the data, without the existence of pragmatic intrusion. Or so we will argue.

Here are the first set of examples that appear to make Levinson’s points, but do not involve apparent temporal readings of “and”:

- (13a) Eating *some* of the cake is better than eating all of it.

<sup>44</sup> If “better-than” expresses a relation between properties, then we can take “Better-than ( $p, q$ )” to be true with respect to a situation if and only if the closest worlds in which  $p$  is instantiated are preferable, in the contextually salient sense, to the closest worlds in which  $q$  is instantiated.

<sup>45</sup> For example, in a context in which health is at issue, that John has spinach is preferable to that John has cake. But in a context in which pleasure is at issue, the opposite may be true.

- (13b) Having a cookie *or* ice cream is better than having a cookie and ice cream.  
 (13c) Having *two* children is better than having three.

We are not certain whether “better than”, in the examples in (13), relates properties or propositions. But let us suppose that they relate propositions.<sup>46</sup> Given the semantics in (12), the examples in (13) all come out true and unproblematic on fairly plausible assumptions. For example, it is quite plausible that, in a given context of use in which the salient preferability relation is health, the most similar worlds in which (say) some cake-eating happens are ones in which some of the cake is eaten, rather than all of it. Similarly, it is plausible that the most similar worlds in which a cookie or ice cream is eaten are ones in which it is not the case that both are eaten. So, given plausible assumptions, these examples do not raise the specter of pragmatic intrusion.

But consider (14):

- (14) Buying *some* of that cake is better than buying *all* of it.

Suppose, in the context in which (14) is uttered, the entire cake is on sale in the supermarket. There is no obvious way to purchase only some of it. In this context, the most similar worlds in which some of the cake is purchased are ones in which all of the cake is purchased. But then the semantic account given in (12) predicts that the truth of (14) requires such worlds to be preferable to themselves (in the contextually determined sense of “preferable”). So (14) should seem radically false. But it seems that (14) can nevertheless be truly uttered.

As Levinson emphasizes, the phenomenon exhibited by (14) occurs with other determiners. For example, suppose that we live in a world in which 90 per cent of all childbirths give rise to triplets. It can nevertheless be true that:

- (15) Having two children is better than having three children.

But, on the assumption that the semantic content of “two” is “at least two”, the semantic clause in (12) appears to entail that (15) cannot be true in this situation. This is counter-intuitive.

There are three general strategies one can exploit in this situation. First, one can retain the semantic clause (12), and defend the thesis that (14) and (15) do not express true propositions relative to the envisaged situations. One might, for example, maintain that utterances of (14) and (15) communicate true propositions, despite

<sup>46</sup> The sentences in (13) are generic sentences. Take (13a); it means something like “Gen(x)(x’s eating some of the cake is better than x’s eating all of it)”. On this account, “better than” relates the proposition expressed by the open sentence “x eating some of the cake” to the proposition expressed by the open sentence “x eating all of the cake”, relative to the assignment function introduced by generic quantifier. So what one is comparing is propositions according to which one person ate some of the cake with propositions according to which that very person ate all of the cake.

the fact that the sentences in context express false ones.<sup>47</sup> Secondly, one can retain the semantic clause in (12), and argue that pragmatic processes affect the intuitive content. The third option is to reject (12), and seek an alternative semantic analysis of a comparative like “better”.

Pragmatic intrusion approaches are versions of the second strategy, the one that Carston and Levinson adopt. In other words, they would retain the semantics in (12), and argue that the scalar implicature of “some” “enriches” the semantic content of (14), and the scalar implicature of “two” enriches the semantic content of (15).<sup>48</sup> On this account, there are strong pragmatic effects on the intuitive content of the utterance, which is the conclusion that Carston and Levinson wish to draw.

However, on both Carston’s and Levinson’s accounts, it is unclear how an implicature “enriches” semantic content. “Enrich” is certainly not a technical term; it is unclear what enriching a semantic content with an implicature is supposed to amount to in the end. In this respect, they have yet to provide an account of the phenomena. As we will now argue, closer attention to the way in which context affects these constructions reveals that only weak pragmatic effects are involved in the cases at hand. That is, we agree with Carston and Levinson that the second strategy is a promising one for accounting for the data. But the way in which pragmatics affects the intuitive content only involves a weak pragmatic effect of context on semantic content, that is, one triggered by the syntax and semantics.

There is a distinctive feature of the cases discussed by these authors, namely that they involve placing focal stress on the relevant expressions. Consider again:

- (13a) Eating *some* of the cake is better than eating all of it.
- (13b) Having a cookie *or* ice cream is better than having a cookie and ice cream.
- (13c) Having *two* children is better than having three.

In each of these cases, one needs to place stress on the italicized expression to obtain the relevant reading.<sup>49</sup> The key to explaining the data lies in recognizing the effects of focal stress.

Giving focal stress to a word makes salient a contextually appropriate contrast class. So, for example, consider:

- (16) John met *Bill*.

<sup>47</sup> See Levinson’s discussions of the “Obstinate Theorist” in ch. 3 of Levinson (2000).

<sup>48</sup> In fact, surprisingly, Carston and Levinson do not provide a semantics for “better than”. Nevertheless, their strategy is clearly to retain a relatively simple semantics for “better than”, and to account for the complexities pragmatically.

<sup>49</sup> Levinson (2000, 400 n. 26) even explicitly suggests placing focal stress on words like “some” in constructions like (13a) if one has trouble with the intuitions.

A speaker may utter (16), with focus on “Bill”, with the intention of asserting that John met Bill, as opposed to another contextually salient person or persons. One way of treating the phenomenon of focus is by saying that placing focus on a constituent gives rise to either the presupposition that the sentence-frame does not hold true of any of the contextually salient alternatives, or the implicature that it does not.<sup>50</sup> Suppose that Frank is the contextually salient person in the context of utterance of (16). Then, on this rough analysis, an utterance of (16) presupposes or implicates that it is not the case that John met Frank.

In the case of (16), a proper name is placed in focus. However, in the examples in (13), determiners (“some”, “two”) and connectives (“or”) are the expressions that are placed into focus. The same general analysis is applicable here, that focusing the expression gives rise to a presupposition or implicature that the sentence-frame is not true of contextually salient alternatives. However, in the case of these expressions, it is not as obvious to the untrained observer what the contextually salient alternatives in fact are.

In fact, determiners such as “some”, and connectives such as “or”, are conventionally associated with contextually salient alternatives, via the “scales” of scalar implicatures. For example, the scale for “some” is  $\langle \text{some, all} \rangle$ , and the scale for “or” is  $\langle \text{or, and} \rangle$ . These scales reflect an intuitive ordering of the “logical strength” of their members.

For example, placing focus on “some” in (17) gives rise to the presupposition or the implicature that the sentence-frame that results from deleting the occurrence of “some” is not true of the other element in the scale associated with “some”:

(17) *Some* bottles are on the table.

In other words, placing focus on “some” in (17) gives rise to the presupposition or the implicature that the sentence-frame “x bottles are on the table” is false as applied to the determiner meaning expressed by “all”. Similarly, placing focal stress on “or” in (18) gives rise to the presupposition or the implicature that the relevant sentence-frame is false for the function denoted by “and”:

(18) John can eat cake *or* John can eat cookies.

So, to sum up, focusing an element that is conventionally associated with a scale affects interpretation by giving rise to a presupposition or an implicature that the sentence-frame is false for the members of that scale that are of greater “strength”.<sup>51</sup>

<sup>50</sup> We are neutral on whether focus gives rise to a presupposition or an implicature in what follows; as we discuss below, in either case, Carston and Levinson’s examples only involve weak pragmatic effects. It is worth mentioning that some theorists (e.g. Glanzberg (2005)) hold that focus is syntactically represented. Since this position would stack the deck in favor of our position, we do not adopt it in what follows.

<sup>51</sup> The connection between focus and scalar implicatures is well documented; see Rooth (1996, 274).

This fact has implications for the evaluation of complex constructions. For example, consider conditionals. On the standard Lewis analysis, a subjunctive conditional is true if and only if its consequent is true in the worlds most similar to the actual world in which its antecedent is true. So, the semantics of subjunctive conditionals, like the semantics of “better than”, involves reference to a similarity relation between worlds. Focusing an element can influence the selection of the similarity relation relevant to the truth-conditions of a particular subjunctive conditional.

Suppose that Frank met Sue at the airport, and invited her to our party, an invitation she declines. Hearing this, John utters:

(19) If *Paul* had met Sue at the airport, she would have accepted the invitation.

John’s utterance is true if and only if the most similar worlds in which Paul had met Sue at the airport and Frank did not meet Sue at the airport are ones in which Sue accepted the invitation. This is so, even if the most likely way for Paul to meet Sue at the airport would be for Paul to go with Frank. Perhaps Paul cannot drive, so, in some more absolute sense, the most similar worlds to the actual world in which the antecedent is true are ones in which Frank and Paul had met Sue at the airport. But such worlds are irrelevant to evaluating the truth of (19) in the envisaged context. By placing focal stress on “Paul” in uttering (19), John forces the selection of a similarity relation that is sensitive to the contrast class.<sup>52</sup>

On one view of focus, focusing a constituent gives rise to a certain implicature or presupposition. If this view is correct, then this implicature or presupposition can affect the intuitive content of a subjunctive conditional. But this is not a strong pragmatic effect. Recall that on the conception of semantics we have adopted, speaker intentions, including those involved in implicatures, can affect semantic content, but only by affecting the choice of referential content in a context of some element in a sentence. In the case of focus and subjunctive conditionals, the way in which the implicature affects the semantic content of subjunctive conditionals fits this model.

We assume that the similarity relation relevant for the semantic content of subjunctive conditionals is traceable to the syntax of the conditional construction, perhaps to the words “if” or “then” themselves. So, the implicature that emerges by focusing a constituent affects the referential content of the element in the syntax of subjunctive conditionals that has, as its referential content relative to a context, a similarity relation. This is a weak pragmatic effect of context on what is communicated, rather than a strong pragmatic effect. Hence, even on the view that focus gives rise

<sup>52</sup> Of course, this point holds for indicative conditionals as well. Suppose we don’t know whether Frank or Paul picked up Sue at the airport, but we do know that if Frank were present, Sue would decline to go to any party. Someone then utters “If *Paul* met Sue at the airport, she is at the party now”.

to an implicature, the fact that focus affects the intuitive content of subjunctive conditionals ends up as merely a weak pragmatic effect.

These facts are directly relevant to the facts cited by Levinson:

(20a) Eating *some* of the cake is better than eating all of it.

(20b) Having a cookie *or* ice cream is better than having a cookie and ice cream.

Focusing “some” in (20a) gives rise to the presupposition or implicature that the sentence-frame “eating *x* of the cake” is false of the other member of the scale associated with “some”, namely the determiner meaning of “all”. This presupposition or implicature does affect the truth-conditions of (20a). But the way it affects the truth-conditions of (20a) is not by “enriching” the semantic content. One dimension of contextual sensitivity of “better than” constructions, according to this semantics, is the relation of similarity between worlds. By focusing “some” in (20a), one invokes a similarity relation according to which the most similar worlds in which some of the cake is eaten are not ones in which all of the cake is eaten. Similarly, by focusing “or” in (20b), one forces a similarity relation in which the closest worlds in which a cookie or ice cream is had are ones in which it is not the case that a cookie and ice cream are had.

In other words, just as with subjunctive conditionals, by focusing the relevant words, one affects the choice of the similarity relation between worlds that is relevant for the truth-conditions of the “better than” construction in that context. So, the truth-conditions of these constructions are affected by scalar facts, but independently of processes such as explicature or implicature “intrusion”. Nor does the scalar information “enrich” the semantic content. Rather, the truth-conditions of “better than” sentences are sensitive to the choice of a similarity relation between worlds, and focus affects the choice of the relation.

To show that the effect here is a weak pragmatic effect, we would have to demonstrate that the similarity relation is the value of some element in the syntax of comparative construction. We have not here provided an explicit syntax for comparatives, and so we cannot justify in detail the thesis that selection of a comparative relation is merely a weak pragmatic effect of context. However, we assume that it is likely that in the final implementation of the semantics in (12), one would trace the introduction of the similarity relation to some element in the expression “better than”. If so, then it is possible to account for all of these alleged examples of pragmatic intrusion without accepting its existence. On the account we have given, one can easily explain the data within the semantics. For, assuming that the similarity relation is the value of some element in the comparative, the account just sketched exploits only weak pragmatic effects in deriving the relevant readings. This is an effect of context on semantic content, given the conception of semantics that we adopted in the first section.

The semantics given in (12) is no doubt oversimplified. And, as mentioned, we have not produced an analysis of the syntax of these constructions to justify our claim that

the similarity relation is the value of some element in the syntactic structure of these constructions. Our point has merely been to show that it is not difficult to describe an account of the phenomena discussed by Carston and Levinson that does not invoke pragmatic intrusion, and indeed only makes appeal to weak pragmatic effects.<sup>53</sup>

Levinson's next set of examples involves indicative conditionals (see Carston (1988, sect. 7)). Here are some of his examples:

- (21a) If you ate *some* of the cookies and no one else ate any, then there must still be some left.
- (21b) If the chair *sometimes* comes to department meetings that is not enough; he should always come.

Levinson's point is that the intuitive content of (21a) is (22a), and the intuitive content of (21b) is (22b):

- (22a) If you ate some but not all of the cookies and no one else ate any, then there must still be some left.
- (22b) If the chair sometimes but not always comes to department meetings that is not enough; he should always come.

But if the intuitive content of the sentences in (21) is as given in (22), then it would seem that the implicatures of the antecedents affect the intuitive content of the sentences in (21).

However, Levinson is wrong to claim that these readings of the sentences in (21) are due to pragmatic intrusion. Rather, the readings in (22) of the sentences in (21) are genuinely semantic, due again to the focal stress placed on the words "some" and "sometimes"; Levinson's pervasive use of italics throughout his examples itself strongly suggests that this is how Levinson means them to be read (see again Levinson (2000, 400 n. 26)). Without placing stress on "some" in (21a), it is not naturally read as having the reading (22a). But this suggests that the relevant readings of the examples in (21) comes from the interaction between the semantics for focus and the semantics of the indicative conditional, rather than the purported phenomenon of pragmatic intrusion.

We assume a simple semantics for conditionals, where both indicative and subjunctive conditionals receive the same analysis in terms of possible worlds, the difference between them being due to the similarity relation relevant for their truth-conditions.<sup>54</sup> As Robert Stalnaker has argued, indicative conditionals normally exploit a similarity

<sup>53</sup> Levinson (2000, 201) also gives examples of what he claims are cases in which manner implicatures affect the intuitive content of sentences containing "better than". However, we find these examples straightforwardly unconvincing, and so do not discuss them here.

<sup>54</sup> Thanks to Brian Weatherson for discussion of the following treatment of conditionals.

relation that counts only those non-actual worlds compatible with the mutually accepted background assumptions as similar worlds for purposes of semantic evaluation. In contrast, subjunctive conditionals involve a similarity relation that reaches outside the worlds compatible with the mutually accepted background assumptions. An indicative conditional is true if and only if the consequent is true in every one of the most relevantly similar worlds in which the antecedent is true. With Stalnaker, we assume that the actual world is always the most similar world, so that non-actual worlds are only relevant for the semantic evaluation of an indicative conditional when the antecedent is false (Stalnaker (1999c, 69)). Finally, we again assume the syntax (and its semantics) of the conditional triggers a search for the contextually relevant similarity relation.<sup>55</sup>

Placing the above rough remarks about focus together with this straightforward analysis of the indicative conditional, we may analyze Carston and Levinson's examples as follows. The focal stress on "sometimes" and "some" gives rise to a presupposition or implicature. For example, in (21a) the speaker intends to give rise to the presupposition or the implicature that not all of the cookies have been eaten by the hearer; this corresponds to the scalar implicature of "some". Similarly, in (21b) the speaker intends to give rise to the presupposition or the implicature that the chair does not always come to department meetings. So doing affects the choice of the contextually salient similarity relation for the sentences in (21). For example, in (21a), by focusing "some", the speaker forces a similarity relation in which all the most similar (non-actual) worlds are ones in which the speaker did not eat all the cookies. This account directly yields the prediction that the conditionals in (21), with "some" and "sometimes" focused, have the truth-conditions given by the sentences in (22).

Let us go through this reasoning in detail with one example. Consider again:

(21a) If you ate *some* of the cookies and no one else ate any, then there must still be some left.

By focusing "some", the speaker forces the selection of a similarity relation for the evaluation of (21a) according to which the most similar (non-actual) worlds are ones in which "you ate *x* of the cookies" is false of the member of the scale associated with "some", which is the determiner meaning of "all".<sup>56</sup> In other words, by focusing "some", the speaker selects a similarity relation in which the most similar worlds do not include any worlds in which the addressee ate all the cookies. So, in the most

<sup>55</sup> One difference between this treatment of indicative conditionals and Stalnaker's treatment is that we do not adopt Stalnaker's selection function, which singles out a unique closest world for the semantic evaluation of a conditional even when the antecedent is false. Fans of conditional excluded middle may retain Stalnaker's assumption if they like.

<sup>56</sup> Henceforth, we will suppress "non-actual".



similar worlds in the context set, it is not the case that the hearer ate all of the cookies. Applying the simple semantics for the indicative conditional to (21a) with this similarity relation, one then considers all of the most similar worlds in which the hearer ate some of the cookies and no one else ate any. Since all of these worlds are already worlds in which the hearer did not eat all of the cookies, the consequent is true in all such worlds.

It therefore falls out from an independently motivated account of focus, and an independently motivated account of indicative conditionals, that the sentences in (21) have the same truth-conditional effect as the sentences in (22). This effect is achieved without any appeal to pragmatic intrusion.<sup>57</sup> Of course, as with the discussion of “better than”, definitively establishing our case that these examples involve only weak pragmatic effects would require demonstrating conclusively that the similarity relation for indicative and subjunctive conditionals is the value of some element in their syntax (e.g., called upon by the lexical meaning of “if”). We believe this to be plausible, but do not claim to have made the case for it here.

In the preceding discussion we have emphasized the role of *focus* in the evaluation of conditionals. In particular, we have shown that the “alternative set” for a focused element in a conditional may affect the semantic evaluation of that conditional, by affecting the choice of the similarity relation for that conditional. We take this to be an instance of the more general fact that the similarity relation for a conditional is affected by the assumptions made by conversational participants. As we discuss below, there are examples similar to the ones discussed by Carston and Levinson, in which no element is focused. Nevertheless, the same kind of response is available for such examples. For the general point is the familiar one that the similarity relation for indicative and subjunctive conditionals is determined as a function in part of the psychological states of the participants in conversational contexts. If there is an element in the syntax that is assigned a similarity relation, relative to a context, such examples are simply more evidence of weak pragmatic effects.

Levinson gives several other alleged examples of kinds of pragmatic intrusion in conditionals. The first kind of example is as in:

- (23) If you have a baby and get married, then the baby is strictly speaking illegitimate. (Levinson (2000: 206 (25a)))

Here Levinson again relies upon the thesis that the semantic content of a conjunctive sentence is not sensitive to temporal information. Given this assumption, the semantic

<sup>57</sup> If the effect of focus is to give rise to a presupposition, the account of the data in (21) is even smoother. For in this case, the focus-induced presuppositions will by definition be part of the context set for evaluation of the indicative conditional, and no appeal to a focus-induced shift of a similarity relation is needed.

content of the conjunct does not entail that the birth precedes the marriage. Then, the information that the first conjunct temporally precedes the second will have to be an implicature, one that furthermore enriches the semantic content of the conditional. However, we have already seen above that the assumption behind this argument, that the semantic content of a conjunctive sentence is not sensitive to temporal information, is incorrect. So these examples of Levinson (2000 (25a–d)) bear no additional comment.<sup>58</sup>

The rest of Levinson's examples also rely on controversial assumptions about what is and is not part of semantic content proper. For example, consider Levinson's example (26a) (2000, 206):

- (24) If Bill and Penny drive to Chicago, they can discuss sociolinguistics in the car for hours.

The antecedent of this conditional may have the intuitive content that Bill and Penny drive to Chicago together. Levinson assumes that the information that Bill and Penny are driving together to Chicago is not part of the semantic content, but rather is what he calls a "together" implicature. However, most semanticists working on collective and distributive readings of verbs would certainly not treat this as a pragmatic implicature but rather as part of semantic content. For example, the semantic content of the antecedent of (24), on some theories, would be taken to be a proposition concerning a driving event with a plural agent.<sup>59</sup>

A further class of examples, again made salient by Robyn Carston (1988), involve causal relations. Consider:

- (25) Mr Jones has been insulted and he's going to resign. (Carston, 1998)

This sentence is naturally understood as communicating that there is a causal connection between the fact that Jones was insulted and his resignation. However,

<sup>58</sup> Another group of Levinson's examples (2000, 207 (26c) and (26c')) rely on the thesis that the semantic content of possessive constructions, such as "Bill's book is good", is not sensitive to the contextually salient possession relation. This claim is not supported by the sort of detailed discussion of the syntax and semantics of possessive constructions that would be required to support it, and so is not worthy of lengthy consideration.

<sup>59</sup> An anonymous referee for this chapter commented that if the collective–distributive distinction is due to a genuine ambiguity that affects semantic interpretation, as we suggest here, this would have the false consequence that any utterance of the sentence "If Bill and Penny drive to Chicago, Bill will get there first (as his car is the fastest)" would be semantically anomalous. But this objection is confused. If there is a genuine ambiguity between collective and distributive readings, then an utterance of the envisaged sentence would not be semantically anomalous, since the person who utters the sentence would intend the distributive reading of the antecedent of the conditional (according to which there are driving events with different agents).

it would not be in the spirit of Grice's "modified Occam's Razor" to postulate a distinctive causal sense of "and".<sup>60</sup> After all, consider:

(26) John took out his key and opened the door.

A sentence such as (26) is naturally understood as communicating that John's door-opening was a result of the fact that John took out his key. If one postulated a distinctive sense of "and" to account for the apparent causal reading of (25), one may have to postulate another distinctive sense of "and" to account for the apparent resultative reading of (26). Such sense-multiplication is to be avoided if possible.

However, one can account for Carston's data without postulating distinctive causal or resultative senses of "and". As Jennifer Saul (2002) has suggested, there are general (and we think defeasible) pragmatic rules governing the order of listing of reported events (rules that can, as we have seen, affect the semantic evaluation of the temporal element associated with verbs). In particular, the maxim of manner requires speakers to "list events in the order in which they occurred, and to cite causes before effects" (Saul (2002, 362)). As a result, unless they believe that the maxim of manner is being flouted, interpreters generally conclude from utterances of (25) that the first event causes the second, and in the case of (26), that the second is a result of the first.<sup>61</sup>

However, there is a more complex group of sentences, concerning which, contra Saul, no such account is plausible.<sup>62</sup> These cases involve conjunctions embedded inside conditionals, such as:

(27) If Hannah insulted Joe and Joe resigned, then Hannah is in trouble.

(27) seems to express the proposition that if Hannah insulted Joe and Joe resigned as a result of Hannah's insult, then Hannah is in trouble. Saul (2002, 363) claims that there is a "perfectly reasonable [Gricean] explanation" of these facts as well, one that "will precisely parallel" the explanation of utterances of sentences such as (25) and (26). However, Saul overstates matters when she suggests that an explanation in terms of the maxim of manner of the intuitive truth-condition for utterances of (27) is "precisely parallel" to the explanation of utterances of sentences such as (25). There are significant differences between the two explanations, which make an explanation in terms of manner in the case of utterances of sentences such as (27) considerably less plausible.<sup>63</sup>

<sup>60</sup> This principle reads "Senses [of words] are not to be multiplied beyond necessity". For Grice's discussion of it, see Grice (1989b, 47 ff.).

<sup>61</sup> Saul herself assimilates this phenomenon to *implicature*, stating that the speaker implicates that the second event is a cause or result of the first. We ourselves are agnostic as to whether this is best described as an implicature.

<sup>62</sup> We are particularly grateful to Robyn Carston for forcing us to treat this style of example in detail.

<sup>63</sup> Saul is aware of these differences (2002, 363). But she disagrees that they impinge on the plausibility of the explanation. It is also worth mentioning that the example Saul herself discusses is a case of

In the case of Saul's explanation of utterances of (25) and (26), the speaker expresses a true proposition, and, by saliently adhering to the maxim of manner, communicates a proposition that is informationally richer. It is perfectly plausible, in this case, to take the speaker as intending to express the true proposition that does not involve causal information, and thereby communicating an informationally richer proposition that does. However, the purportedly parallel account of utterances of (27) is different. In this case, speakers would typically express a false proposition not involving causal information, and thereby communicate a true proposition that does. Furthermore, in this latter case, it is not particularly plausible to suppose that speakers are aware of the false proposition they express in uttering (27).<sup>64</sup> These are significant differences between the two explanations.

Of course, speakers do sometimes knowingly express false propositions, and thereby communicate true ones. This is what happens, for example, in non-literal speech. But we are not considering cases of non-literal speech here; someone who assertively utters (27) does not intend to be speaking non-literally. There should be a strong presumption against treating speech that is intended literally (and used correctly) on the model of non-literal speech. In sum, the purported explanation for the intuitions in (27) has significant theoretical costs, ones that the explanation for the intuitions in (25) and (26) do not have. These costs should be unattractive whatever one's theoretical commitments. However, for those who reject semantic skepticism and semantic modesty, they are simply untenable.

Fortunately, however, there is no reason to give a non-semantic account of the intuitive readings of (27). The relevant reading of (27) is simply predicted by the semantics for indicative conditionals that we have endorsed. In a context in which the speaker has in mind a causal relationship between Hannah's insulting of Joe and Joe's resignation, all relevantly similar worlds in the speaker's context set in which Hannah insulted Joe and Joe resigned will *ipso facto* be ones in which Joe's resignation is due to Hannah's insult. The speaker's context set is what is epistemically open to her. This may include worlds in which the conjunction holds, and there is no causal relationship between the conjuncts. But given that it is salient that she has a causal relationship in mind, such worlds will not be the most relevantly similar worlds in

apparently temporal "and" embedded in the antecedent of a conditional. As we have made clear above, such temporal readings are generated *within* the semantics, though of course not by an alleged temporal reading of "and". There is no need for a post-semantic explanation of such cases.

<sup>64</sup> It is worth noting that the explanation here is even more extreme than pragmatic accounts of quantifier domain restriction. Defenders of such accounts can argue that the false proposition they claim to be expressed by someone who utters a sentence like "Everyone is at the party" is easily accessible ("Do you really mean *everyone*?"). However, the false proposition supposedly expressed by (27), on this pragmatic account, is not so easily accessible. It is therefore harder to see how the speaker could intend to say it, or intend to make as if to say it.

the context set. So, if she has a causal relation in mind between the two events, that is just to say that the similarity relation for indicative conditionals will select those worlds in which there is a causal relationship between the conjuncts of the antecedent as the most similar worlds to the world of utterance in which the antecedent is true. So, the causal reading of (27) is predicted by the simple semantics for the indicative conditional that we have adopted above.

We have argued in this section that the examples of pragmatic intrusion given by Carston and Levinson are not convincing. The intuitions behind these examples can be accommodated within semantic interpretation, without requiring strong pragmatic effects to explain them. We think a general moral follows from this investigation. Before claiming that a set of intuitions cannot be due to semantic interpretation, theorists need to have investigated all of the semantic options. For, as Levinson admirably acknowledges, claims about what can only be derived pragmatically may very well be vitiated by subsequent syntactic and semantic investigation.

## Conclusion

An important obligation of the philosophy of language is to provide a clear distinction between semantics and pragmatics. However, at the center of the philosophy of language and cognitive science there has been a sustained debate about the scope and interest of semantic content, with many theorists arguing that semantic content plays a marginal role in an explanation of linguistic behavior. These debates have been clouded by disagreement over the proper definitions of “semantic” and “pragmatic”. Our intention in this chapter has been to provide a clear characterization of semantic content, and then use it to evaluate the debate about the scope and interest of semantic content.

As we have emphasized, these different positions on the scope of semantic content are, at this stage of research, merely educated guesses. Many philosophers and cognitive scientists maintain that, once syntactic and semantic inquiry are finished, it will turn out that the semantic content of a sentence relative to a context is not a good guide to what speakers typically use that sentence to communicate. This position has radical consequences for standard methodology in philosophy. If it is true, it is unclear, for example, whether any of the strategies canvassed at the beginning of this chapter are legitimate. In the final section we have addressed one of the strongest arguments in favor of skepticism about the scope of semantic content, and shown it to be unpersuasive. Our view is that, as yet, such skepticism remains unwarranted.

## References

- BACH, KENT (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BACH, KENT (2002a), 'Semantic, Pragmatic', in J. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 284–92.
- BACH, KENT (2002b), 'Seemingly Semantic Intuitions', in J. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 21–33.
- BRAUN, DAVID (1993), 'Empty Names', *Noûs*, 27/4: 449–69.
- BRAUN, DAVID (1994), 'Structured Characters and Complex Demonstratives', *Philosophical Studies*, 74: 193–21.
- BRAUN, DAVID (1996), 'Demonstratives and their Linguistic Meanings', *Noûs*, 30/2: 145–73.
- CAPPELEN, HERMAN (2002), 'Insensitive Quantifiers', in J. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 197–213.
- CAPPELEN, HERMAN, and ERNIE LEPORE (1997), 'On an Alleged Connection between Indirect Speech and the Theory of Meaning', *Mind and Language*, 12/3 and 4: 278–96.
- CARSTON, ROBYN (1988), 'Implicature, Explicature, and Truth-Theoretic Semantics', in R. Kempson (ed.), *Mental Representations: The Interface between Language and Reality* (Cambridge: Cambridge University Press), 155–81.
- CARSTON, ROBYN (ms), 'Linguistic Underdeterminacy and Pragmatic Enrichment'.
- COHEN, L. J. (1971), 'Some Remarks on Grice's Views about the Logical Particles of Natural Language' in Y. Bar-Hillel (ed.), *Pragmatics of Natural Language* (Dordrecht: Reidel), 50–68.
- CRIMMINS, MARK (1992), *Talk about Beliefs* (Cambridge, Mass.: MIT Press).
- CRIMMINS, MARK, and JOHN PERRY (1989), 'The Prince and the Phonebook: Reporting Puzzling Beliefs', *Journal of Philosophy*, 86/12: 685–711.
- DEROSE, KEITH (1995), 'Solving the Skeptical Problem', *Philosophical Review*, 104: 1–52.
- DEROSE, KEITH (1999), 'Contextualism: An Explanation and Defense', in John Greco and Ernie Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Basil Blackwell), 187–205.
- DEVER, JOSHUA (2001), 'Complex Demonstratives', *Linguistics and Philosophy*, 24/3: 271–330.
- ENC, MURVET (1987), 'Anchoring Conditions for Tense', *Linguistic Inquiry*, 18: 633–57.
- GAUKER, CHRISTOPHER (1997), 'Domain of Discourse', *Mind*, 106: 1–32.
- GLANZBERG, MICHAEL (2005), 'Focus: A Case study on the Semantics–Pragmatics Boundary', in Zoltán Gendler Szabó (ed.), *Semantics versus Pragmatics*, (Oxford: Clarendon Press), 72–110.
- GRICE, PAUL (1989a), 'Logic and Conversation', in Grice, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press), 22–40.
- GRICE, PAUL (1989b), 'Further Notes on Logic and Conversation', in Grice, *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press), 41–57.
- HEAL, JANE (1997), 'Indexical Predicates and their Uses', *Mind*, 106: 619–40.
- HECK, RICHARD (2001), 'Do Demonstratives have Senses?', *Philosophers' Imprint*, 2/2: 1–33.
- HEIM, IRENE (1982), 'The Semantics of Definite and Indefinite Noun Phrases', Ph.D. diss. (University of Massachusetts).

- HIGGINBOTHAM, JAMES (2002), 'Why is Sequence of Tense Obligatory?', in G. Preyer and G. Peter (eds.), *Logical Form and Language* (Oxford: Oxford University Press), 207–27.
- KAPLAN, DAVID (1989), 'Demonstratives', in Perry Almog and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), 481–564.
- KING, JEFFREY (1995), 'Structured Propositions and Complex Predicates', *Noûs*, 29/4: 516–35.
- KING, JEFFREY (2001), *Complex Demonstratives: A Quantificational Account* (Cambridge, Mass.: MIT Press).
- KING, JEFFREY (forthcoming), 'Tense, Modality, and Semantic Value', *Philosophical Perspectives*.
- LEPORE, ERNIE, and KIRK LUDWIG (2000), 'The Semantics and Pragmatics of Complex Demonstratives', *Mind*, 109: 199–240.
- LEVINSON, STEPHEN (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicatures*, (Cambridge, Mass.: MIT Press).
- LEWIS, DAVID (1998), 'Adverbs of Quantification', in Lewis, *Papers in Philosophical Logic* (Cambridge: Cambridge University Press), 5–20.
- LUDLOW, PETER (1995), 'Logical Form and the Hidden Indexical Theory: A Reply to Schiffer', *Journal of Philosophy*, 92: 102–7.
- MONTAGUE, RICHARD (1974), 'Pragmatics', in R. Thomason (ed.), *Formal Philosophy* (New Haven: Yale University Press), 95–118.
- OGIHARA, TOSHIYUKI (1996), *Tense, Attitudes, and Scope* (Dordrecht: Kluwer).
- PARSONS, TERENCE (1990), *Events in the Semantics of English: A Study in Sub-atomic Semantics* (Cambridge, Mass.: MIT Press).
- PARTEE, BARBARA (1984), 'Nominal and Temporal Anaphora', *Linguistics and Philosophy*, 7: 243–86.
- PERRY, JOHN (1997), 'Indexicals and Demonstratives', in Robert Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language* (Oxford: Basil Blackwell), 586–612.
- PERRY, JOHN (2001), *Reference and Reflexivity* (Stanford: CSLI).
- RECANATI, F. (1993), *Direct Reference* (Oxford: Basil Blackwell).
- RICHARD, MARK (1983), 'Direct Reference and Ascriptions of Belief', *Journal of Philosophical Logic*, 12: 425–52.
- RICHARD, MARK (1990), *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them* (Cambridge: Cambridge University Press).
- ROOTH, MATS (1994), 'Focus', in Shalom Lappin (ed.), *The Handbook of Contemporary Semantics* (Oxford: Basil Blackwell), 271–97.
- SALMON, NATHAN (1986), *Frege's Puzzle* (Cambridge, Mass.: MIT Press).
- SAUL, JENNIFER (1997), 'Substitution and Simple Sentences', *Analysis*, 57/2: 102–8.
- SAUL, JENNIFER (2002), 'What Is Said and Psychological Reality: Grice's Project and Relevance Theorists' Criticisms', *Linguistics and Philosophy*, 25: 347–72.
- SIEGEL, SUSANNA (2002), 'Perception and Linguistic Demonstration', *Philosophers' Imprint*, 2/1: 1–21.
- SOAMES, SCOTT (1987), 'Direct Reference, Propositional Attitudes and Semantic Content', *Philosophical Topics*, 15: 47–87.
- SOAMES, SCOTT (2005), 'Naming and Asserting', in Zoltán Gendler Szabó (ed.), *Semantics versus Pragmatics*, (Oxford: Clarendon Press), 356–82.
- SPERBER, D., and D. WILSON (1986), *Relevance* (Cambridge, Mass.: Harvard University Press).

- STALNAKER, ROBERT (1999a), *Context and Content* (Oxford: Oxford University Press).
- STALNAKER, ROBERT (1999b), 'Pragmatics', in Stalnaker (1999a), 31–46.
- STALNAKER, ROBERT (1999c), 'Indicative Conditionals', in Stalnaker (1999a), 63–77.
- STANLEY, JASON (2003), 'Context, Interest-Relativity, and the Sorites', *Analysis* 63/4: 269–80.
- STOWELL, TIM (1982), 'The Tense of Infinitives', *Linguistic Inquiry*, 13: 561–70.
- SZABÓ, ZOLTÁN GENDLER (2001), 'Adjectives in Context', in I. Kenesei and R. M. Harnish (eds.), *Perspectives on Semantics, Pragmatics and Discourse* (Amsterdam: John Benjamins), 119–46.
- WETTSTEIN, HOWARD (1984), 'How to Bridge the Gap between Meaning and Reference', *Synthese*, 58: 63–84.
- WILSON, DEIRDRE (1975), *Presuppositions and Non-Truth-Conditional Semantics* (London: Academic Press).



# 5

## Making it Articulated

Here is the view of linguistic communication I find plausible. First, a speaker makes an utterance, and her linguistic intentions uniquely determine a certain syntactic structure, or “logical form”, as it is known in syntax. If her utterance is a successful linguistic assertion, the logical form is sentential. Successful interpretation involves assigning denotations to the constituents of the logical form, and combining them in accord with composition rules that do not vary with extra-linguistic context. The denotations that successful interpreters will assign to constituents of a logical form will be constrained by the linguistic conventions governing those elements. In the case of certain elements, which wear their context-dependent nature on their sleeve, the linguistic conventions governing them are rather lax. These elements are the obvious indexicals, the obvious demonstratives, pronouns, context-dependent quantifiers such as “many” and (perhaps) “that”, and covert pronominal elements whose existence can be demonstrated by purely syntactic tests. What results from a successful application of this first stage of interpretation is a unique proposition, a fully truth-evaluable entity. Furthermore, if the process is successful, the proposition

This chapter originated as a reply to Robyn Carston, initially at the “Mind and Language” conference on pragmatics and cognitive science in Oxford, and then subsequently at the Eastern Division APA. My two greatest debts are to Robyn and to Kent Bach. My interactions with Robyn have taught me a great deal about the motivations and details of the relevance theoretic perspective, and just a great deal about pragmatics generally. I have been discussing the arguments of this chapter at length with Kent for many months, and his numerous e-mails and comments have greatly improved them. The audience at the “Mind and Language” conference was also tremendously helpful; in particular, François Recanatani, Dan Sperber, Rob Stainton, and Deirdre Wilson all made incredibly valuable contributions to my greater edification, and hence to the final result of this chapter. Conversations and e-mail with Michael Nelson also were very helpful.

Richard Heck and Brett Sherman have been urging me for about a year and a half to take more seriously the over-generation worry facing my opponents. In the end, I recognized they were correct, though I have developed this worry differently than they would. So I owe them a significant debt of gratitude. I have talked over every line of this chapter with Jeff King, and he has provided his usual extraordinarily useful suggestions. In addition, discussions with Richard Breheny, Delia Graff, Ernie Lepore, Peter Ludlow, Stephen Neale, Zoltán Gendler Szabó, and Rich Thomason have had substantial positive effect on the final result. Finally, comments by two anonymous referees (and editorial advice by Robyn) resulted in large-scale improvements

it yields is one the interpreter would recognize as the proposition expressed by that assertion. In the final stage of a successful communicative interchange, an interpreter applies standard Gricean reasoning to derive the implicatures of the linguistic act.

Virtually every claim I have made in the description of my project has been severely criticized in recent years, most prominently by Kent Bach, Robyn Carston, François Recanati, Dan Sperber, and Deirdre Wilson. In particular, these theorists hold that, in many cases of successful interpretation, the proposition that the hearer would intuitively believe to be the proposition expressed by the communicative act contains elements that are not the value of any constituent in the sentence uttered, nor introduced by composing these values. Instead, these elements are provided directly by context. Loosely following John Perry, I will call such elements *unarticulated constituents*.<sup>1,2</sup>

If there are unarticulated constituents of utterances, then the result of combining the values assigned to elements of a structured representation is at best an interpretive way-station on the path to the truth-conditions hearers would intuitively associate with the linguistic act. One consequence of relegating the result of semantic interpretation to this role is that intuitions speakers have about what is expressed by linguistic acts are no longer an accurate guide to syntactic structure and semantic content (cf. Bach (forthcoming)). In contrast, I believe such intuitions are an accurate guide to syntactic structure and semantic content (see p. 90). But I recognize that this position needs to be argued for, and not simply assumed.

My purpose in this chapter is therefore to raise doubts about the existence of unarticulated constituents. That is, I want to argue in favor of the view that all the constituents of the propositions hearers would intuitively believe to be expressed by utterances are the result of assigning values to the elements of the sentence uttered, and combining them in accord with its structure. The way I will accomplish this is by questioning the existence of some of the processes that theorists have claimed underlie the provision of unarticulated constituents to the propositions recovered by hearers in linguistic interpretation.

<sup>1</sup> I have heard it said that, for Perry, an unarticulated constituent is one that is not the value of a *pronounced* element. But Perry (1986) argues for the thesis that there are unarticulated constituents of thoughts. If the values of phonologically null representational elements were still unarticulated, then every semantic value of a mental representation would be an unarticulated constituent.

<sup>2</sup> More precisely: an entity (object, property, or function) *e* is an unarticulated constituent relative to an utterance *u* if and only if (a) *e* is a constituent of the proposition that a competent, reflective speaker under normal circumstances would intuitively believe to be what is expressed by *u*, and (b) *e* is not the value of any constituent in the expression uttered in *u*, and (c) *e* is not introduced by context-independent composition rules corresponding to the structural relations between the elements in the expression uttered.

In the first section, I introduce an argument I have exploited in recent work against the existence of unarticulated constituents. In the second section, I discuss an influential paper of Wilfred Sellars, which has led some philosophers of language to the view that there is a legitimate notion of ellipsis that is the process by which unarticulated constituents can be provided. If there is such a process, then it can be exploited to respond to the kind of argument raised in the first section. However, I argue that no evidence has been given by Sellars for the existence of such a process. In the third section, I sketch some strategies theorists have appealed to in defending the existence of unarticulated constituents. Finally, in the fourth section, I argue that the processes that one needs to postulate to evade the arguments against unarticulated constituents described in section two lead to unacceptable consequences.

## I The Argument from Binding

In recent work, I have given the following sort of argument against the existence of unarticulated constituents. First, I consider a construction utterances of which allegedly involve unarticulated constituents. I then show that there are readings of the construction in question that require recognizing that a variable exists in the syntactic structure of the relevant construction, whose value, relative to a context, is the allegedly unarticulated constituent. If so, then the constituent is not unarticulated after all.

Here is an example of the argument at issue. An utterance of (1) can communicate the proposition that every bottle that John just bought is in the fridge:

- (1) Every bottle is in the fridge.

That is, an utterance of (1) can communicate a proposition about a restricted domain of bottles. For example, an utterance of (1) can communicate the proposition that every bottle Sally bought on 11 April 2001 is in the fridge. Let us call this phenomenon *quantifier domain restriction*.

On one natural account of the phenomenon of quantifier domain restriction, the syntactic structure of (1) does not contain a variable whose value, relative to a context, is a restricted domain. Rather, denotation assignment and composition yields the proposition that every bottle (unrestricted) is in the fridge. Successful interpretation then involves augmenting this minimal proposition with material from extra-linguistic context.<sup>3</sup> Call this the pragmatic account.

<sup>3</sup> The proposition is “minimal” because, although it may be involved in the derivation of the proposition the hearer would believe to be expressed in the speech act, it is not one the hearer would necessarily recognize as the proposition thus expressed.

The pragmatic account of quantifier domain restriction can be challenged with the following sort of argument, which I will call *the argument from binding*. Consider a sentence such as:

- (2) Every student answered every question.

An utterance of this sentence may communicate the proposition:

- (3) Every student answered every question *on her exam*.

These facts challenge the pragmatic account; and this in two related ways. First, these facts provide evidence for a covert pronominal element, the content of which, in this context, is the same as that of the italicized material in (3). For positing such an element, and holding it to be bound by the quantifier “every student”, delivers the relevant reading (more on this below). Secondly, these facts challenge the pragmatic account, because it is not clear how to capture the bound readings of such sentences given the resources available to the advocate of the pragmatic account. Let me explain each of these points in turn.

Much syntactic structure is unpronounced, but no less real for being unpronounced. This fact raises the question of how to detect unpronounced syntactic structure. A working hypothesis is that an unpronounced element exists in the structure of a sentence just in case there is behavior that would be easily explicable on the assumption that it is there, and difficult to explain otherwise. For example, here is a classical argument that the agent of a passive is actually still syntactically present, though unpronounced. In control theory, the theory of the nature of unpronounced elements like “PRO” in English (the subject of infinitival clauses), there is strong evidence that “PRO” must have a syntactically local controller. This, after all, is what explains the ungrammaticality of:

- (4) \*The ship sunk to collect the insurance.  
 (5) \*The record broke in winning the race.

(4) and (5) are ungrammatical, because their syntactic structure is:

- (6) The ship sunk [PRO to collect the insurance]  
 (7) The record broke in [PRO winning the race]

And these occurrences of “PRO” have no potential local controller; or, rather, the available local controllers (“the ship” and “the record”) are not expressions that denote things capable of collecting insurance or winning races. However, (8) and (9) are perfectly in order:

- (8) The ship was sunk to collect the insurance.  
 (9) The record was broken in winning the race.

This strongly suggests that the agent by-phrase in a passive is in fact syntactically present after all. This is a standard way to argue that a constituent that is not phonologically realized is nevertheless syntactically present.

The argument from binding is an argument of the very same structure. One characteristic syntactic feature of pronouns is their capacity to be bound by variable-binding operators. By demonstrating the existence of bound readings of quantifier-domain variables, one provides evidence of behavior that is explicable on the assumption that there is an unpronounced pronominal element, and difficult to explain otherwise. This raises the second challenge to the pragmatic account, which is to explain the bound readings of the sentences in question without postulating a variable in the syntactic structure of quantified noun phrases. My main purpose in this chapter is to evaluate various attempts to respond to this second challenge.

It is worth emphasizing why it is important to focus on this second challenge to the pragmatic account. Some of those who reject my account of linguistic communication accept many syntactic arguments for unpronounced linguistic structure of the sort just discussed (cf. Bach (2001, 37–8)). However, they are virtually uniform in their rejection of the argument from binding. Presumably, the reason theorists such as Bach, who, despite the similarities between the argument from binding and the argument concerning passives, reject the former but not the latter is that they believe that there are viable pragmatic accounts of the data concerning binding, but not of the data concerning passive constructions. If there is no tenable pragmatic account of the data concerning binding, then this position is refuted.

So the question I want to address in the remainder of this chapter is whether or not one can account for the data discussed in examples (2) and (3) without postulating a variable whose value, relative to a context, is a quantifier domain.<sup>4</sup> In the next section, I turn to a critique of an influential discussion by Wilfred Sellars, which has led some philosophers to accept the existence of an interpretive strategy that could be marshaled to respond to the argument from binding. In the following sections, I consider more recent suggestions.

<sup>4</sup> There is a complication here, because on the account of quantifier domain restriction I find most promising (see pp. 95–103), quantifier domain restriction is not just due to the presence of one variable in the syntactic structure of quantified sentences, but rather two variables. For example, on this account, the reason (2), relative to a context, can express the same proposition that (3) expresses, relative to the same context, is that the common noun “question” is associated with two variables, a function variable and an object variable. The object variable is bound by the initial quantifier “every student”, and the function variable receives as its value the salient function from individuals to quantifier domains, which in this case is a function taking individuals to the property of being something on that individual’s exam.

## II Sellars on “Ellipsis”

In Wilfred Sellars famous paper “Presupposing” (Sellars, 1954), Sellars makes a distinction between what he calls three “types of ambiguity” (ibid., 198ff.). The first type of ambiguity is ordinary lexical ambiguity. The second type of “ambiguity” discussed by Sellars is the “ambiguity” associated with “egocentric particulars”, or what we would now call demonstrative and indexical expressions, namely that these expressions can change their reference as a function of the context of use.<sup>5</sup> The third type of “ambiguity” discussed by Sellars is introduced via the following example, involving two different interchanges between two mathematicians:

Jones: Seven is divisible by three.

Smith: Seven is not divisible by three.

Jones: Seven is.

(later) Jones: Seven is divisible by four.

Smith: Seven is not divisible by four.

Jones: Seven is.

As Sellars writes: “each of these dialogues contains an utterance of ‘Seven is’; and it is clear that what is communicated by these utterances is a function of the contexts in which they are uttered” (ibid., 200). Sellars then proceeds to argue that this third kind of “ambiguity” is genuinely distinct from the kind of “ambiguity” that is involved in the interpretation of demonstrative and indexical expressions. According to Sellars, “Correctly made utterances of [‘This is red’] are *complete* even though they say what they do by virtue of their context. On the other hand, the two utterances of ‘Seven is’ are as such not complete and are only made complete by the context in which they are uttered” (ibid.). Sellars then introduces the term “ellipsis” as a name for the kind of ambiguity that is present in Jones’s two utterances of “Seven is”.

Sellars then proceeds to argue that sentences containing “incomplete” definite descriptions, such as “The table is large”, suffer from this third kind of ambiguity, namely what he calls “ellipsis”.<sup>6</sup> That is, Sellars claims that a sentence such as “The

<sup>5</sup> I will use the term “indexical” broadly, to apply to any context-sensitive expression.

<sup>6</sup> An “incomplete definite description” is a descriptive phrase, the nominal component of which is true of more than one thing. The problem such phrases raise for a Russellian account of definite descriptions is that, on a Russellian account, definite descriptions only denote if the nominal component of the description is true of exactly one thing. But it seems that atomic sentences containing incomplete definite descriptions (such as “The table is large”) can, relative to some contexts, express truths, *contra* to what the Russellian theory appears to predict. As Neale (1990, 93) has clearly explained, the problem of “incomplete” definite descriptions is simply a species of the more general problem of quantifier domain restriction.

table is large” is incomplete in *just the same way* as Jones’s utterances of “seven is”.<sup>7</sup> According to Sellars, an utterance of “The table is large” is elliptical for a sentence like “The table over there is large” in just the same way as Jones’s first utterance of “Seven is” is elliptical for “Seven is not divisible by three”.

Sellars’s claim here is startling, as is his use of the term “ellipsis” to describe the phenomenon of incomplete definite descriptions. With characteristic incisiveness, Peter Strawson, in his reply to Sellars in the very same journal issue (Strawson (1954)), gives a clear explanation of the oddity of both Sellars’s claim and his vocabulary. Since the point is important, I quote Strawson at length:

In Section II of his article, Sellars distinguishes between (among others) two kinds of “ambiguity”. The first is manifested by sentences containing certain words, among them the word “this”. Sentences such as “This hat is yours”, “This wine is good” . . . may be uttered in many different situations, and in each case the context or setting of the utterance will be an essential element in the determination of the reference made by the use of the phrases “this hat”, “this wine”, etc. Contrasted with these are sentences which manifest a different kind of “ambiguity”, and which Sellars calls “incomplete” or “elliptical”. Examples analogous with his own would be the sentence “James is,” which might be uttered as an answer to the question “Who is going to drive?” or the question “Who is going to walk?”; or the sentence “Castor oil is,” which might be uttered as a rejoinder to the assertion “Castor oil isn’t harmful” or to the assertion “Castor oil isn’t horrible”. Now it is clear that there are many differences between the two classes of sentences here contrasted by Sellars, and many differences between his two sorts of ambiguity. It also seems very reasonable to call the second set of sentences, as opposed to the first, incomplete or elliptical. If one had to justify these phrases, I think one would be inclined to say that the sentences were *formally, linguistically* deficient, that they did not come up to a certain standard of how a non-conversational English sentence should be composed; and one would point out that in their conversational setting, the deficiency is remedied by the *linguistic* context, that the surrounding *remarks* supply the missing words.

But Sellars’ next suggestion I find utterly puzzling. For he says that such a sentence as “The table is large” is incomplete or elliptical in the same sense as sentences of his second class; that *this* sentence has *this* kind of ambiguity. I fail to see any reason whatever for saying this . . . Of course there are differences between “the” and “this”. But there are also close resemblances between the ways in which context, in the widest sense, helps to determine the reference of many “the” phrases and the ways in which context helps to determine the reference of many “this” phrases. And there are no ways peculiar to the former, as opposed to the latter, in which their contextual dependence resembles the contextual dependence of incomplete or elliptical sentences. (Strawson (*ibid.*, 222–4)).

<sup>7</sup> Actually, Sellars maintains that sentences containing “incomplete” definite descriptions suffer from both the second and third kind of ambiguities, since such sentences are said to be elliptical for sentences containing demonstrative expressions. But this point is irrelevant here.

Strawson makes several points in this passage. First, he states that Sellars is justified in using the term “elliptical” for utterances of sentences like “seven is” and “James is”, in contexts in which the surrounding linguistic context “supply the missing words” to the incomplete sentences, because, in the ordinary sense of the term “elliptical”, an utterance is elliptical just in case its surrounding linguistic context quite literally supplies the missing words. This point is about vocabulary, the proper use of the term “ellipsis”. Secondly, Strawson emphasizes that Sellars is incorrect to claim that the very same phenomenon is occurring with incomplete definite descriptions as with elliptical utterances of “seven is” or “James is”. Third, Strawson emphasizes that the way context works to resolve the “incompleteness” of definite descriptions is considerably more similar to the way context works to resolve the reference of demonstrative expressions than it is to the way context works to resolve elliptical utterances. In other words, the kind of context dependency at issue with “The table is large” is much more like the context-dependency of “This table is large” than it is the context-dependency of “seven is” in the discourse between Jones and Smith.

All of Strawson’s points are correct. First, the terms “elliptical” and “incomplete” only seem apt when the phenomenon involves words that are supplied to incomplete sentences by linguistic context.<sup>8</sup> Secondly, and more importantly, once one distinguishes between the kind of processes operative in resolving the reference of occurrences of demonstratives, and the kind of processes operative in resolving elliptical sentences (such as Sellars’s example “seven is” and Strawson’s example “James is”) it seems bizarre to claim, as Sellars does, that the processes at work in resolving the context-dependence of “The table is large” are the latter rather than the former.

There is a sense, however, in which Strawson has won the battle against Sellars, only to lose the war. Strawson is right, as against Sellars, that the processes at work in resolving the context-dependency of definite descriptions are very similar to the processes at work in resolving the context-dependency of demonstrative and indexical expressions. But this provides little comfort to Strawson’s thesis that definite descriptions are themselves referring expressions. For the right theory of quantifier domain restriction can explain Strawson’s insight, without giving up the thesis that definite descriptions are quantificational expressions.

According to the account of quantifier domain restriction I favor, the reason that “The table is large” may be used to express a truth is that the noun “table” is

<sup>8</sup> I emphasize Strawson’s first point only because Neale (2000) claims that Sellars’s use of “elliptical” is the “ordinary, non-technical sense” of the term (*ibid.*, 287), which according to Neale does not bring any commitment whatsoever to context supplying missing words to a sentence. Strawson’s first point is that the term “ellipsis” only seems appropriate to describe Sellars first set of examples of what he calls ellipsis, which are clearly cases in which linguistic context supplies missing words to the sentence uttered.



associated syntactically with domain indices, which are assigned values by context in a similar manner as demonstrative expressions are assigned values by context. The phenomenon of “incomplete” definite descriptions therefore turns out to be simply a species of indexicality, as Strawson correctly claimed it was. One can therefore accept Strawson’s insight that the contextual processes at work in the resolution of indexicality are the same as those at work in the resolution of “incomplete” definite descriptions, without endorsing his view that definite descriptions themselves are in any way indexical or demonstrative expressions. In fact, Strawson’s insight provides an *additional argument* for accounts of quantifier domain restriction that treat it as a species of indexicality.<sup>9</sup>

Be that as it may, Sellars’s discussion has engendered much confusion in subsequent literature on context-dependence. Sellars’s paper has been given new life thanks to the influential discussion of “incomplete” definite descriptions in Neale (1990, 95 ff.), from which much work on the topic in the subsequent decade borrowed its vocabulary. Neale’s discussion takes off from an uncritical acceptance of Sellars’s vocabulary and distinctions. But Sellars’s vocabulary and distinctions suffer from ambiguities, unclarity, and errors, which, because of Neale’s discussion, continue to disrupt the literature on the topic today.<sup>10</sup>

To begin with, there are three distinct ways of taking Sellars’s discussion, all of which are consistent with the text of Sellars’s article. According to the first, Sellars is suggesting that in both his examples “seven is” and “The table is large”, context supplies actual missing words, which are not pronounced in the utterance.<sup>11</sup> According to the second way of understanding Sellars’s suggestion that there is a common explanation of both his examples “seven is” and “the table is large”, *neither* case involves the supplementation of hidden words by context (to adopt Strawson’s more natural usage of the term “ellipsis”, on this reading of Sellars, neither type of case is a case of ellipsis). Finally, the third way of taking Sellars’s discussion is that, though the first kind of example involves the supplementation of hidden words by context, and the second case does not, there is nevertheless some common interpretive genus of which these are species.

<sup>9</sup> It is worth mentioning that on the most complete and thorough existing account of complex demonstratives, namely King (2001), complex demonstratives are quantificational expressions like definite descriptions, whose nominal complements also are associated with domains by context (*ibid.*, 128 ff.).

<sup>10</sup> In particular, every paper on “incomplete” definite descriptions that uncritically adopts Neale’s vocabulary of “explicit” and “implicit” approaches to context-dependence must be examined to see whether it is vitiated by Neale’s reliance on Sellars’s discussion.

<sup>11</sup> So on this interpretation, when Sellars writes (1954, 200), “a given utterance of [‘The table is large’] is elliptical and states what would be nonelliptically stated, for example, by ‘The table over here is large’”, what Sellars means is that the sounds made in the utterance event of “The table is large” reflect the syntactic structure of “The table over here is large”.

On the first and second interpretation of Sellars's discussion, what he would be claiming is generally acknowledged to be false. The first interpretation is generally acknowledged to be false, because it yields an implausible account of the context-dependency of sentences like "The table is large". The second interpretation is generally acknowledged to be false because it yields an implausible account of examples such as Sellars's "seven is". That leaves us with the third interpretation. On this interpretation, although Sellars's example of "seven is" is a case of genuine ellipsis, and his example of an utterance of "The table is large" is not, there is a common interpretive genus of which these are species. Let us call this technical notion introduced by this interpretation of Sellars, *magical ellipsis*.

Suppose there is a process of magical ellipsis, which hearers may exploit, and which speakers may exploit, knowing (perhaps only tacitly) that hearers exploit it in interpretation. Then one could respond to the argument from binding straightforwardly. One could simply maintain that a speaker can utter (2), and successfully communicate what would have been expressed by an utterance of (3), because (2) is magically elliptical for (3).

But this is not a satisfactory response to the argument from binding. First of all, no reason has been given to believe in that there is a common interpretive genus process at work both in resolving the "incompleteness" of quantified expressions and in resolving the incompleteness of the utterances of "seven is" in the scenarios Sellars describes. But there is a more important point. It is that a simple appeal to a process like magical ellipsis is not an *explanation*. An utterance of (2) can communicate the same proposition as would have been expressed by an utterance of (3). This is the fact that needs explanation. Appealing to magical ellipsis amounts to answering this question by stating that there exists some interpretive process that allows a shorter sentence to communicate the same proposition as would have been expressed by an utterance of a sentence that contains additional words. This is akin to "explaining" the perturbations caused in the orbit of Uranus that actually led to the postulation of Neptune by stipulating that there exists an astronomical process that allows a planet to move as if there were another planet causing those perturbations.<sup>12</sup>

What is required to respond to instances of the argument from binding is a detailed explanation of the nature of an interpretive process that would allow a speaker to utter (2), and thereby communicate what would have been expressed by an utterance of (3). This is what theorists such as Bach, Carston, Recanati, Sperber, and Wilson have attempted to provide. In the next section, I turn to a discussion of some of the processes suggested by these theorists, and how they could be exploited to respond to the argument from binding.

<sup>12</sup> Similarly, what Neale (1990, 2000) calls the "explicit" approach to context-dependence, simply amounts to a re-description of the phenomenon to be explained, rather than an account of it.

### III Pragmatic Binding?

In this section, I sketch a process that could underlie the provision of unarticulated constituents. The process I will describe is similar to the ones sketched in Bach (2000) and Carston (forthcoming), which they use to respond to the argument from binding discussed in section II. However, since Bach and Carston have rather different theoretical commitments, I will not keep to the letter of their discussions.<sup>13</sup> In all relevant respects, however, the process I describe for the provision of unarticulated constituents is similar to the processes they endorse in responding to the argument from binding for the existence of variables.

Let us suppose, for the moment, that when a speaker utters a sentence containing a predicative use of a comparative adjective, as in the sentences in (10), the sentence uttered by the speaker does not have a syntactic position whose value, relative to a context, is a comparison class:

- (10) a. John is tall.  
 b. That lamp is cheap.

It is uncontroversial that a speaker can, with the use of such sentences, communicate propositions that concern particular comparison classes (say, 10-year-old boys, for an utterance of (10a)). So, if the supposition is right, the comparison class must enter into interpretation in some manner other than being assigned to an element of the sentence uttered. That is, if the supposition is right, then comparison classes for predicative uses of comparative adjectives are unarticulated constituents of the propositions communicated by utterances of sentences containing such constructions.

Let us continue to suppose that when a speaker utters a sentence containing a predicative use of a comparative adjective, as in the sentences in (10), there is no syntactic position in the sentence whose value is the comparison class. If this assumption is correct, then the semantic content of sentences such as (10a) and (10b), relative to a context, is not a proposition. Rather it is what Kent Bach calls “a propositional radical” (Bach (1994, 127)). In these cases, the semantic content would be a function from comparison classes to propositions. But somehow, by uttering sentences whose semantic contents are propositional radicals, speakers manage to communicate genuine propositions. How could this occur?

<sup>13</sup> The most important difference between Bach and theorists such as Carston, Recanati, Sperber and Wilson, and Charles Travis is that Bach holds that there is a notion of what is said that is strictly the result of semantic interpretation, and plays a privileged role in an account of communication. However, according to Bach, what is said in this sense is often not a proposition, and, even when so, often not the proposition the hearer would intuitively recognize as what is expressed by utterance. So this difference is irrelevant for our purposes.

Suppose that interpreters, when they hear a sentence whose semantic content in context is either less than fully propositional, or expresses a proposition that the speaker clearly does not intend to communicate, quite automatically replace the uttered sentence by another one, one that contains additional words. This more complicated sentence, relative to that context, would semantically express the proposition the speaker intends to communicate, and the interpreter in fact grasps. But the recovery of this more developed sentence, although often fairly automatic and almost always unconscious (or tacit), is still a pragmatic process. Perhaps the interpreter tacitly recognizes that the speaker was being less than fully explicit, and hence that an ostensible violation of a Gricean maxim has occurred. This ostensible violation of a Gricean maxim then guides the replacement of the less explicit sentence by the more explicit one.<sup>14</sup> Or perhaps the speaker's utterance triggers a pragmatic process by which the interpreter tacitly sorts through a set of propositional representations for the most relevant among them, as in the relevance theoretic framework of Sperber and Wilson (1986).

If interpreters tacitly exploit these interpretive processes, then speakers will be tacitly aware that hearers may exploit these sorts of processes. Speakers may then use their tacit knowledge that interpreters employ these processes to utter sentences whose semantic content, relative to a context, is less than fully propositional, or is a proposition that is quite distinct from the one they wish to communicate. Speakers can engage in this communicative sloth, because they know that interpreters will use tacit pragmatic reasoning to obtain the correct proposition. In this way, a regular process of semantically expressing one content, and successfully communicating a very different one, can be sustained, even though the content the interpreter would believe to be expressed by the communicative act was in fact recovered by a pragmatic process of enriching the original sentence.

Suppose some story like this is true. Now consider an application of the argument from binding, this time to a predicative use of a comparative adjective (see p. 55):

(11) Most species have members that are small.

An utterance of (11) can communicate the proposition that most species *S* are such that *S* has members that are small for *S*. If it is normal for speakers to exploit a tacit interpretive process that replaces a sentence by a lengthier one, then perhaps this is what occurs when speakers utter (11). Perhaps (11) does not, after all, contain a comparison class variable bound by the quantifier "most species". Then, the semantic content of (11) is not a proposition, but a propositional radical. But speakers may utter (11) to communicate the proposition that most species *S* are such that *S* has members that are small for *S*, by exploiting the tacit pragmatic interpretive process

<sup>14</sup> Thanks to Kent Bach for discussion here.

that interpreters use when speakers are less than fully explicit. In this case, interpreters tacitly recognize that the speaker has not semantically expressed what she wished to communicate, and replace sentence (11) by (12), which does semantically express what the speaker wished to communicate:

(12) Most species have members that are small for that species.

Here, the interpretive process at play is no different from the one at play in supplying unarticulated constituents to the propositions communicated by ordinary utterances of sentences containing predicative uses of comparative adjectives.<sup>15</sup>

So if the story told in this section is correct, pragmatic processes can supply variables. Speakers are tacitly aware that their interlocutors have at their disposal an interpretive strategy that allows them to recover the proposition the speaker intends to communicate, even in cases in which this proposition is not semantically expressed by the words used. Speakers use this knowledge in saving breath, in uttering sentences that do not express propositions, or do not express the propositions they wish to communicate. Since the pragmatic interpretive process is tacit, hearers are not consciously aware that they exploit it. Hence, normal intuitions about what is expressed are not sensitive to what is in fact semantically expressed. For if the story told in this section is true, a tacit pragmatic process often intervenes between grasp of what is semantically expressed and grasp of what is communicated, even in cases in which the hearers are not consciously aware that the speaker has intended to communicate something different than she expressed.

#### IV Too Much of a Good Thing?

In the last section, I described a pragmatic interpretive strategy, the existence of which would explain how speakers could communicate propositions expressed by sentences with quantifier-variable interactions, despite the fact that the sentences they use to communicate them do not syntactically involve variables. In this section, I reject the existence of such processes, on the grounds that they are not consistent with everyday communicative limitations. I conclude by extending these arguments more generally to a large class of “free enrichment” pragmatic processes.

On the hypothesis we are considering, speakers tacitly know that, if they utter a sentence, which, for lack of a variable or pronominal expression, does not semantically

<sup>15</sup> An exactly parallel explanation would be given for how an utterance of (2) can communicate what would have been expressed by an utterance of (3), except in this case the semantic content of (2) would be propositional.

express a proposition, or fails to express the proposition they wish to communicate, interpreters will exploit a tacit pragmatic process that results in the replacement of the uttered sentence by another sentence that contains the missing pronominal element. Furthermore, interpreters will do this tacitly, without realizing that they are engaging in this pragmatic process. Thus, the process of supplying missing variables or pronouns is no different from the process of providing unarticulated constituents generally (on this story).

Here is the first reason to think that this hypothesis is implausible. It is a long-standing presumption in syntax that vacuous quantification results in particularly bad grammaticality violations. However, if the hypothesis we are now considering were true, it would be mysterious why vacuous quantification leads to ungrammaticality. For example, consider:

(13) Everyone<sub>j</sub> who John ran, he liked.

If there were the sort of pragmatic process that allowed an interpreter to replace smoothly an uttered sentence by another sentence that contains a new bound variable, the ungrammaticality of (13) would be a complete mystery. For all an interpreter has to do to make (13) grammatical is (tacitly) add on the phrase “by  $x_j$ ” to the sentence, yielding:

(14) Everyone<sub>j</sub> who John ran by  $x_j$ , he liked.

And of course (14) is perfectly grammatical. Similarly, consider:

(15) Everyone has had the privilege of having John greet.

(15) is wildly ungrammatical. But if the hypothesis under consideration were true, then the ungrammaticality of (15) would be a mystery. For all an interpreter has to do to make (15) grammatical is to add the little word “her”, yielding:

(16) Everyone has had the privilege of having John greet her.

If there were the sort of pragmatic process of the sort endorsed by theorists such as Bach and Carston, why could it allow interpreters to smoothly add “on her exam” to (2), but not the considerably shorter “her” to (15)?

Here is the problem posed by these sentences for the defender of our hypothesis. If bound variables can be provided by pragmatics to allow interpreters to grasp propositions that they otherwise would not have, why wouldn't those same tacit pragmatic processes be at work in making utterances of ungrammatical sentences appear to be utterances of grammatical ones? If our semantic intuitions are unwittingly responsive to the semantic features of a sentence that is the result of a tacit pragmatic process, then one would naturally expect our syntactic intuitions to be responsive to sentences that are the result of such processes as well. In short, if there are tacit

pragmatic processes that undermine our apparently *semantic* intuitions, it is a mystery why they wouldn't also undermine our apparently *syntactic* intuitions.<sup>16</sup>

Here are several replies to this worry. The first involves denying that pragmatic processes can intervene in syntactic processes. The problem with this claim is that it is inconsistent with the use of free enrichment to provide variables to be bound, given that binding is a syntactic process. Furthermore, some defenders of free enrichment argue that pragmatics enters into even the most straightforwardly syntactic mechanisms, such as verb phrase ellipsis. The worry raised here is of course of *particular* concern for such theorists.

Here is a second reply to the worry. In the relevance theoretic framework championed by Robyn Carston, Dan Sperber, and Deirdre Wilson one may stipulate that the *cost* of repairing an ungrammatical sentence is so high, that to do so by pragmatic means is never acceptable. However, the notion of cost here is mysterious. Even if one's life depended upon it, it would still be impossible to hear (13) and (14) as grammatical. So, "cost" is a technical term. We currently lack a sufficiently independent grip on the notion of the "cost" of reinterpreting a sentence to evaluate the claim that the cost of repairing (13) and (14) is too high. So this notion must be explicated in such a way that it is a consequence of how it is explicated that the cost of repairing an ungrammatical sentence is far higher than giving it an interpretation it does not have. Without such an explanation, the claim that the cost of repairing an ungrammatical sentence by pragmatic means is always too high remains a mere stipulation.<sup>17</sup>

Here is a final reply to the worry.<sup>18</sup> One might respond that the tacit pragmatic processes under consideration are ones that only are operative when the sentence uttered expresses a complete proposition. Since (13) and (15) do not express propositions, interpreters do not apply the tacit pragmatic processes. However, this response is a non-starter. To respond to the argument from binding, the relevant pragmatic processes need to be operative, even when the sentence uttered does not express a proposition. For example, on the account discussed in the last section, (11) (repeated here) does not express a proposition in any context:

(11) Most species have members that are small.

<sup>16</sup> This worry is not as significant for a purely semantic notion of free enrichment, where pragmatic processes operate to supply constituents directly to the result of semantic interpretation, without the hearer having to go through the step of replacing the uttered sentence by a more developed sentence. But such a notion of free enrichment is particularly ill-suited to deal with the argument from binding, since to treat such cases, context would have to supply denotations for bound variables (see Stanley and Szabó (2000)).

<sup>17</sup> Thanks to Dan Sperber for discussion that led to this paragraph.

<sup>18</sup> Thanks to an anonymous referee for this objection.

Yet the same tacit pragmatic processes are supposedly in effect to provide a variable to the expanded sentence. So the advocate of our hypothesis cannot accept this restriction on the relevant tacit pragmatic processes, because otherwise she would have to accept that the argument from binding successfully establishes the existence of variables for cases such as (11).<sup>19</sup>

But let us suppose that the defender of these tacit pragmatic processes can reply to this worry. In fact, let us allow her the straightforward stipulation that pragmatic processes can only operate on grammatical sentences. Even allowing this, (13) and (15) raise a second, entirely distinct, worry for the defender of the hypothesis.

Many ungrammatical sentences are nevertheless, in context, interpretable. That is, very often, an utterance of an ungrammatical sentence can nevertheless be corrected or repaired by pragmatic means. In contrast, (13) and (15), even in context, are extremely difficult to interpret. On the hypothesis we are considering, interpreters can grasp smoothly propositions that would be expressed by sentences containing bound variables when presented with utterances of sentences that do not contain those variables. If this hypothesis were correct, then interpreters should be able to exploit this very pragmatic device to correct or repair utterances of (13) and (15). Thus, if the hypothesis under consideration were correct, then, despite the ungrammaticality of (13) and (15), they should nevertheless be easily interpretable in context. The tacit pragmatic process described in the previous section should naturally produce the interpretations in (14) and (16), respectively. But (13) and (15) are well-nigh impossible to interpret. This fact is mysterious, if there are tacit pragmatic processes that can allow interpreters to grasp propositions expressed by sentences with quantifier-variable structures, when presented with utterances of sentences whose syntactic structures lack such structure.

This leads to a third worry for the defender of the hypothesis. For if there were the sort of pragmatic processes postulated by the hypothesis, then there are countless examples of sentences that speakers should be able to use, without consciously and obviously flouting Gricean maxims, to communicate propositions that they in fact cannot thereby communicate.

Here is one example. Suppose that speakers tacitly know that hearers have available to them a tacit pragmatic process that allows them to grasp, when they hear a sentence that does not semantically express what the speaker clearly intends to communicate, the proposition that would be expressed by a sentence that results from the uttered sentence by the addition of some extra syntax, including bound pronouns. If so, then

<sup>19</sup> In other words, in the vocabulary of Recanati (1993), appealing to a distinction between distinct primary pragmatic processes of “saturation” and “free enrichment” does not help in the least to respond to this objection, because both processes will be implicated in the provision of variables.



speakers should not only be able to exploit these processes tacitly in communicating what would be expressed by (3) by an utterance of (2), but they should also be able to exploit them by smoothly communicating what would be expressed by any of the sentences in (18) by an utterance of (17):

- (17) Everyone likes Sally.  
 (18) a. Everyone likes Sally *and himself*.  
 b. Everyone likes Sally *and his mother*.

In short, if there were the envisaged pragmatic processes, then speakers would be able to utter (17), and thereby successfully communicate (say) (18b), via the audience's exploitation of the alleged tacit pragmatic process.<sup>20</sup> But this is not possible. Therefore, the hypothesis that there are such processes stands refuted.

One reply to this argument is to maintain that there are in fact contexts in which (17) can be used to communicate what utterances of (18a) or (18b) would express.<sup>21</sup> It is uncontroversial that our background knowledge about exams helps us in grasping the information that every student answered every question on her exam from an utterance of:

- (2) Every student answered every question.

According to one way of advancing this reply, when (17) is uttered in a context in which it is commonly known that everyone likes his mother, it in fact communicates what would have been expressed by (18b).

Suppose we are at a party in which only those acquaintances of mine who like and appreciate their mothers have been invited. In a discussion with my friend Hannah, we start to discuss whether we should invite our mutual charming yet eccentric friend Sally over at the last minute. Hannah worries that there are people at the party who are offended by Sally's ebullience. To assuage her, I utter (17). Despite the fact that both Hannah and I are fully aware that everyone at the party likes his mother, my utterance of (17) clearly does not communicate what would have been expressed by (18b). For example, the information that everyone likes his mother is certainly no part of my speaker's meaning in uttering (17).

So, there is a distinction between information that is part of the background knowledge in a conversation, and information that is communicated in a speech act. It is obvious that (17) can be uttered in a context in which it is part of the background knowledge that everyone in the context likes his mother. What is not obvious is that

<sup>20</sup> In fact, the process in question would even be an instance of the somewhat restricted pragmatic process Recanati (1993, 261 f.) calls "strengthening", since the output proposition of such a process would entail the input proposition.

<sup>21</sup> Thanks to Michael Nelson for pressing this point forcefully.

(17) can be used to communicate what (18b) would express, except via the exploitation of obvious violations of Gricean maxims. Though it is obvious that there are contexts in which it is known that everyone in the situation likes his mother, simple appeal to such contexts is of no help at all in making the case that (17) could be used to communicate what (18b) would express in that context.

Now, there are in fact cases in which (17) could be used to communicate what an utterance of (18b) would have expressed. Suppose that you and I both know that Sally has such incredibly unique motherly qualities, so that it is well known that everyone who likes Sally also likes his mother. With this in mind, I utter (17), perhaps while arching my eyebrow knowingly, thereby communicating that everyone here likes his mother. This is a case in which (17) can be used to communicate what an utterance of (18b) would have expressed in that context.

However, this case also is of no help to someone who wants to show that (17) can be used to communicate what would be expressed by (18b), via the exploitation of tacit pragmatic processes. For in this case, the proposition that everyone here likes his mother is a straightforward fully conscious Gricean implicature of the utterance. A reflective, competent speaker would recognize that what was expressed by the utterance of (17) was the proposition that everyone at the party likes Sally. What the defender of the existence of the sort of tacit pragmatic process discussed in the last section requires is an example in which an utterance of (17) successfully communicates what would be expressed by (18b), relative to that context, without there being any obvious conscious violation of a Gricean maxim to communicate an implicature. If there are such tacit pragmatic processes, then there must be such cases, since speakers, tacitly knowing that interpreters have recourse to such an interpretive strategy, would exploit this knowledge in using (17) to communicate (18b). But it is not possible to use (17) to communicate what would have been expressed by (18b), except by obviously flouting a Gricean maxim.<sup>22</sup> It follows that there are not the sort of pragmatic processes described in the last section.

One might respond to the arguments presented in this section by accepting that there are serious problems with the existence of tacit pragmatic processes that allow sentences without pronouns to communicate bound readings, but deny that this raises a problem for all tacit pragmatic strategies that allow an interpreter to replace one sentence by a more explicit one. However, if Bach (2000) and Carston (forthcoming) are correct, this position is not tenable. If there are tacit pragmatic strategies that allow an interpreter to grasp the proposition expressed by a more explicit sentence upon hearing an utterance of a less explicit version of that sentence, it is difficult to see how these processes could be constrained to disallow interpreters from replacing the less

<sup>22</sup> In the vocabulary of Kent Bach, though it is possible to utter (17) and thereby *implicate* that everyone likes his mother, it is not possible to utter (17) with (18b) as the *implicature*.

explicit sentences by ones containing phrases with bound pronouns. As we have seen, it appears unlikely that there are these latter pragmatic processes. This sheds doubt on the existence of any pragmatic processes that allow a speaker tacitly to replace an uttered sentence by a more explicit one. This poses a significant worry for theorists as diverse as Bach, Carston, Neale, Recanati, Sperber, and Wilson, all of whom centrally rely on the existence of processes of these sorts.

## References

- BACH, K. (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BACH, K. (2000), 'Quantification, Qualification, and Context: A Reply to Stanley and Szabó', *Mind and Language*, 15: 262–83.
- BACH, K. (2001), 'You Don't Say?', *Synthese*, 128: 15–44.
- BACH, K. (forthcoming), 'Seemingly Semantic Intuitions', in Keim-Campbell, O'Rourke, and Shier (eds.), *Topics in Contemporary Philosophy: Meaning and Truth* (New York: Seven Bridges Press).
- CARSTON, R. (forthcoming), 'Explicature and Semantics', in S. Davis and B. Gillon (eds.), *Semantics: A Reader* (Oxford: Oxford University Press).
- KING, J. (2001), *Complex Demonstratives* (Cambridge: MIT Press).
- NEALE, S. (1990), *Descriptions* (Cambridge: MIT Press).
- NEALE, S. (2000), 'On Being Explicit: Comments on Stanley and Szabó, and on Bach', *Mind and Language*, 15: 284–94.
- PERRY, J. (1986), 'Thought without Representation', *Supplementary Proceedings of the Aristotelian Society*, 60: 137–52.
- RECANATI, F. (1993), *Direct Reference* (Oxford: Blackwell).
- SELLARS, W. (1954), 'Presupposing', *Philosophical Review*, 63: 197–215.
- SPERBER, D., and WILSON, D. (1986), *Relevance* (Cambridge: Harvard University Press).
- STANLEY, J., and SZABÓ, Z. G. (2000), 'Reply to Bach and Neale', *Mind and Language*, 15: 261–5.
- STRAWSON, P. (1954), 'A Reply to Mr. Sellars', *Philosophical Review*, 63: 216–31.

# 6

## Semantics in Context

Consider an utterance of the sentence “Some philosophers are from New York”. If no philosopher in the world comes from New York, competent speakers of English know that it is false. They also know that this utterance is true, if six philosophers in the world come from New York. In other words, competent English speakers have clear intuitions about the conditions under which what is said by an utterance of this sentence is true or false.

The apparent *source* of such intuitions is not difficult to locate. Competent English speakers know the meanings of the words in the sentence “Some philosophers are from New York”. They also know how to combine the meanings of each of the words in this sentence to arrive at what is said by the utterance of the sentence, “Some philosophers are from New York”. It is that linguistic competence that seems to be the source of their ability to report correctly about the truth of what is said by that sentence relative to different possible circumstances, for example, the circumstance in which there are no philosophers from New York, or the circumstance in which six philosophers come from New York.

So, the explanation for our ability to report about the truth and falsity of what is said by an utterance of “Some philosophers are from New York” in various possible situations is as follows. Competent English speakers know the meanings of the words used, and understand how they are combined. Their grasp of the truth-conditions of the utterance of that sentence is due to their ability to combine the meanings of the words, relative to the context of utterance.

With this explanation in mind, consider an utterance of the sentence “Every philosopher is from New York”, made at a small philosophy conference. It is natural to take this utterance to say something that is true if and only if every philosopher *at the conference* is from New York. If we cleave to the model of understanding just described, we will seek to explain our understanding of the truth-conditions of this utterance by appeal to a process of combining the elements of the sentence “Every philosopher is from New York”, using our understanding of the words used in the sentence. But, of course, there appears to be no expression in the sentence “Every philosopher is from New York” that corresponds to the understood constituent expressed by “at this conference”.

Similarly, suppose, pointing at a 5-foot tall, 7-year-old child, I utter the sentence “He is tall”. I am most naturally understood as saying something that is true if and only if the child in question is tall *for a 7-year-old child*. Preserving the model of understanding we began with, according to which our intuitions about the truth-conditions of an utterance are due to a process of combining meanings of the parts of the sentence uttered, would require us to find some constituent in the sentence that could be taken to supply the understood property of *being a 7-year-old child*. But again, it appears that the sentence “He is tall” contains no such constituent.

So, we have a predicament. If we look at certain sentences, there seems to be a clear and elegant explanation of why we have the intuitions we do about the truth conditions of utterances of those sentences. But if we consider utterances of other sentences, the explanation appears to break down. The first response to this predicament is to attempt to preserve the clear and elegant explanation in the face of the apparently recalcitrant data. The second is to abandon the clear and elegant explanation of the source of our truth-conditional intuitions in favor of a different one.

My concern with the second response to the predicament is that the suggestions I am aware of for dealing with the additional complexity essentially end up abandoning the project of giving a systematic explanation of the source of our intuitions. They invariably involve appeal to unconstrained and non-explanatory notions or processes. I have therefore been inclined to pursue the first of these options.

My purpose in this chapter is to continue the project of defending the clear and elegant explanation of the source of our intuitions about the truth-conditions of utterances. I will do so by considering some replies to previous arguments in favor of it. I will argue that proponents of abandoning the clear and elegant explanation have not yet made their case.

## I The Challenge from Context-Sensitivity

On the simple explanation of the source of our intuitions about the truth-conditions of utterances of sentences we understand, it is due primarily to a compositional process of interpretation. Our knowledge of meaning, together with our knowledge of relevant contextual facts, allows us to assign meanings to the parts of a sentence, and the intuitive truth-conditions of an utterance of that sentence are what results from combining these values. Somewhat tendentially, I will call proponents of the simple explanation, *semanticists*.

Innumerable researchers from pragmatics have challenged the semanticist’s model. Here is the form of the standard challenge. First, a linguistic construction C is produced that appears intuitively to have a certain reading R. Secondly, the researcher claims

that the readings cannot be due to the semantics of that construction. That is, the claim is that R cannot be due to the compositional semantic interpretation of C, relative to the envisaged context of use. The conclusion the researcher draws is that the assumption that the intuitive truth-conditions of a sentence relative to a context are due to semantics is incorrect.

A large number of researchers opposed to the semanticist employ arguments of this sort (a brief list of the most prominent exponents includes Kent Bach, Herman Cappelen, Robyn Carston, Ernie Lepore, Stephen Levinson, Francois Recanati, Dan Sperber, Charles Travis, and Deirdre Wilson; there are many more). Typically, such researchers do not just supply a single example, but a list of disparate examples. For example, the following is a representative list that could occur in any one of a hundred papers written in the past decade by researchers in this tradition:

- (1) John is tall. (for a fifth grader)
- (2) John is finished. (with grading)
- (3) Every boy (in the class) is seated.
- (4) John and Mary went to Paris (together/separately)
- (5) If Lincoln hadn't gone to the theater, he wouldn't have been assassinated.  
(fixing certain background assumptions)
- (6) John ate breakfast. (this morning)
- (7) John had breakfast this morning. (in the normal way, through his mouth)
- (8) John ate. (mushrooms)
- (9) The ham sandwich (person who ordered the ham sandwich) is getting annoyed.
- (10) The apple is green. (on the inside)

The bracketed material is intended to indicate the material that cannot be provided by semantics, but only by pragmatics.

As one may imagine, any such list will include cases that virtually all proponents of the simple model of interpretation believe are uncontroversially generated by the compositional interpretation of the sentence uttered. For example, there is much recent investigation into the syntax and semantics of gradable adjectives that generates the supposedly pragmatic material as the semantic value of some element in the syntax (either a comparison class variable, or a degree variable). Semantic treatments of plurals treat collective-distributive ambiguities in a variety of ways. For example, some treat collective-distributive ambiguities as structural ambiguities due to the relative scope of an event quantifier. On such a treatment, the two readings of (4) are due to scope facts in the syntax, rather than pragmatics. Since Partee (1973), most linguists have defended the view that verbs are associated with temporal variables, that have their references filled via deixis. On this view, generating the relevant reading of (6) simply involves speaker intentions determining the value of a temporal variable in

the syntactic structure of the sentence uttered. Finally, few semanticists would balk at associating the provision of an accessibility relation for counterfactuals to some element in the syntax, either the conditionals words themselves, or some covert element.

On the other side of the spectrum, some of the examples that are provided in such lists seem to require pragmatic treatment, on the grounds that the alleged intuitive truth-conditions are richer than those delivered by tutored intuitions about truth-conditions. For example, as I will argue in the final section, with the help of recent work by Luisa Marti, it is clearly accessible to a native speaker of English that it is no part of the truth-conditions of an utterance of (7) that John ingested his breakfast through his mouth. Of course, when someone tells us that John ate this morning, we assume he did so in the normal way. But no one would deem an utterance of (7) *false* if, contrary to default assumptions, they discovered that John ingested breakfast in some non-standard way, such as being spoon fed. So the manner of eating is no part of the intuitive truth-conditions of (7), but is rather pragmatically conveyed information. We also assume other things when we hear an utterance of (7), for example, that John's breakfast wasn't prepared by a Martian. But none of this is information that is carried semantically, and, *pace* Carston (2002, 203) and Wilson and Sperber (2002), it is odd to suppose that anyone has ever advanced a theoretical position that would commit them otherwise.

Nevertheless, between the two extremes I have just discussed, there are some examples that are genuinely worrying for the semanticist. For instance, it certainly appears that the intuitive truth-conditions of an utterance of (9) involve a person, rather than a ham sandwich. Yet it is not clear that a process that maps ham sandwiches onto persons counts as genuinely semantic. To take another example, the intuitive truth-conditions of (3) certainly involve reference to a domain of quantification. But if domain restriction is a matter of information being freely provided by context, that too does not seem to be a process that can be considered genuinely semantic.

Some have tried to respond to this predicament by arguing that the semantic content of a sentence, relative to a context, is only a minimal part of the intuitive truth-conditions of that utterance, a version of what King and Stanley (Chapter 4 of this book) call *semantic modesty*. As King and Stanley emphasize, the worry with this response is that it is unclear what role the minimal semantic core ends up playing in an account of the intuitive truth-conditions, if one accepts that processes such as free enrichment account for much of our intuitions in examples such as (1)–(10). If free pragmatic enrichment is a process hearers regularly use to interpret utterances, and speakers are aware of this, then why can't speakers utter sentences whose semantic content is minimal or vacuous, and rely on such pragmatic processes to do the bulk of the expressive work? So I am not sanguine about semantic modesty as an intermediate position for the semanticist.

So much the worse for the semanticist, one might think. However, if our intuitions about truth and falsity are responsive to processes that are not linguistically controlled, we need an *explanatory* account of information freely provided by context. And it is not clear that such an account is in the offing. The most serious problem facing the advocate of free pragmatic enrichment to intuitive truth-conditions is that of *over-generation* (cf. Stanley (Chapter 5 of this book)). If our intuitions about the truth-conditions of utterances of quantified sentences are due to a process of free pragmatic enrichment, then it would be a mystery why utterances of certain sentences lack certain readings. For example, why is it the case that an utterance of (11) can express the same proposition as an utterance of (12), but never the same proposition as an utterance of (13)?

- (11) Every Frenchman is seated.
- (12) Every Frenchman in the classroom is seated.
- (13) Every Frenchman or Dutchman is seated.

In short, since these “enrichment” processes are not linguistically constrained, they should be constrained only by general pragmatic reasoning. But why do general pragmatic facts allow (11) to express (12) but not (13)?

In the light of these worries with free pragmatic enrichment accounts of intuitive truth-conditions, it is important to investigate the possibility that the intuitive truth-conditions of utterances of sentences such as those in (1)–(10) are due to linguistically determined content. Linguistically determined content is content that is constrained not just by pragmatic means. Particularized conversational implicatures, for example, are constrained only by pragmatic means, and hence are not part of linguistically determined content. In contrast, the value of a term such as “she”, relative to a context, is linguistically determined, because the speaker intentions that determine its value must be referential intentions consistent with the literal meaning of “she”. So the question posed by such examples is how to establish that the intuitive readings of the problematic sentences in (1)–(10) are due to linguistically determined content.

## II Responding to the Challenge

I adopt the conception of semantics at work in Chapter 1 and spelled out in detail in Chapter 4 of this book. The semantic content of a sentence relative to a context is derived by taking the semantic contents of the parts of that sentence, relative to that context, and composing them in accord with the composition rules governing the syntactic structure of that sentence. The semantic value of a basic constituent



of a sentence is what is determined by speaker intentions together with features of the context, in accord with the standing meaning of that lexical item. Given this conception of semantics, the position of the semanticist is then that the source of our intuitions about the truth and falsity of utterances relative to various possible circumstances is due to semantics.

When faced with the claim that a certain construction *C* has a reading *R* that *prima facie* does not seem traceable to the semantics, the semanticist has three options. The first option is to establish that the alleged reading is not part of the intuitive truth-conditions of an utterance of that sentence, but is instead due to the pragmatics (as in the discussion of (7), above). The second option is to argue that the claim that reading *R* is not due to the semantics is due to an overly simplistic conception of the semantic content of some elements of *C*. When the correct semantics for the relevant expression is given, the reading does emerge from the semantics (cf. King and Stanley (Chapter 4, sect. V), on “implicature intrusion”). The third option is to argue that the claim that reading *R* is not to the semantics is due to an overly simplistic conception of the syntactic structure of *C*. In fact, *C* contains covert structure, and once this is recognized, reading *R* does emerge from the semantics (see Chapter 2 on domain restriction, and Chapter 1 for discussion of other constructions).

So, when faced with a list such as that given in the previous section, the semanticist has, in each case, three choices. The first is to reject the semantic significance of the data; the second is to give an alternative semantic assignment to some overt element; and the third is to argue for covert syntactic structure. As I have indicated, it is a construction specific matter which of these options is preferable. The difficulty facing the semanticist’s opponent is that she must establish, for each case, that none of the three very different alternatives is available as an account of the data.<sup>1</sup>

Of all the constructions on the list, I think the central worry for the semanticist is (9), the case of deferred reference. Not only is there a strong intuition that the deferred meaning is part of the intuitive truth-conditions, but the deferred meaning enters into certain linguistic processes, such as anaphora and ellipsis. For example, the natural reading of (14) is one in which the anaphoric element “his” receives its value not from the “literal” content of “the ham sandwich”, but from its deferred meaning:

(14) The ham sandwich wants his bill now.

Similarly, when we consider someone uttering (15) in the kitchen of a restaurant, describing the predicaments of two waiters, it is the deferred meaning of “an annoying ham sandwich” that is carried over to the ellided constituent:

(15) Bill served a ham sandwich, and John did too.

<sup>1</sup> The only pragmaticist I know of who seems to recognize the daunting challenge this poses to the opponent of the semanticist is Stephen Levinson (cf. Levinson (2000, 214)).

In particular, (15) cannot be interpreted as conveying that Bill served a person who ordered a ham sandwich, whereas John served a ham sandwich. Finally, one could argue that the literal meaning of an expression provides a guide for its deferred meaning, and so the deferred meaning is semantic after all.

I think neither of these points show that deferred reference is semantic. In a nominative metaphor such as (16a), we see the same phenomenon as in (14), and in (16b), we see the same phenomenon as in (15), where the metaphorical reading is what is carried over in ellipsis:

- (16) a. The pig in the next room wants his check immediately.  
 b. John is a pig, and Bill is too.

We cannot interpret (16b) to mean that John is a person who is a sloppy eater, and Bill is (for example) a pig, perhaps John's pet. But on a standard view of metaphor, metaphor is not semantic. If the metaphorical meaning of an expression does not affect the semantic content of sentences containing it, relative to a context, then the fact that deferred reference behaves in a similar manner should not lead us to believe that deferred reference is semantic.

Joseph Stern (2000, 69–70) has recently used facts such as (16b) to argue that metaphor affects semantic content, that (as he would put it), there is such a thing as semantically significant metaphorical *meaning*. But, as Elizabeth Camp (ms) has pointed out, we also see the same phenomenon with *irony*. Consider:

- (17) John: Bill is a fine friend.  
 Sally: Sue is too.

If John's utterance is intended ironically, then the ellided constituent "fine friend" in Sally's utterance must be understood ironically as well. But this does not show that irony is semantic, or that there is such a thing as semantically significant ironical *meaning*. As Stern (*ibid.*, 232) writes, "Now, whatever controversy surrounds the status of metaphorical meaning, the ironic 'meaning' of an utterance is surely not a semantic meaning." So, such ellipsis facts do not demonstrate that a phenomenon is semantic.

The second argument that deferred reference is semantic is that the literal meaning of an expression in context provides a guide to its deferred meaning. For example, the literal meaning of "the ham sandwich" provides a guide to the deferred meaning of "the ham sandwich" in (14), which is *the person who ordered the ham sandwich*. So if the mark of the semantic is guidance (in some sense) by literal meaning, then there is evidence that deferred reference is semantic.

But it is also the case that the literal meaning of "has nice handwriting", in the context of an utterance of "John has nice handwriting" in a reference letter, provides a guide for the implicated property, *is a bad philosopher*. So the fact that the literal meaning

is used in deriving the deferred meaning does not show that the deferred meaning is linguistically controlled in the relevant sense (i.e., semantic).<sup>2</sup>

I think a more general argument can be given that deferred reference should not be treated as semantic. The mark of the semantic is that semantic content is *constrained* by linguistic meaning. At the very least, the semantic content of an expression, relative to a context, must be something of which that expression is true. If it is not, it is hard to see how the semantic content of that expression has been constrained by the conventional meaning of that expression. But in the case of deferred reference, that is not true. If deferred reference were semantic, the denotation of “the ham sandwich” would be something of which the predicate “ham sandwich” were not true. So it is hard to see how the deferred reference of “the ham sandwich”, in (9), is semantic, since it is not constrained by the conventional meaning of the words used. So, one theoretical consideration that should lead us to deny that deferred reference (as in example (9)) is not semantic is that the deferred reference of an expression is not something of which the conventional meaning need be true.<sup>3</sup>

A second consideration involves the *scope* of the phenomenon. One reason against taking metaphor to be semantic is that virtually any term can be used metaphorically. This suggests that metaphor has to do with the *use* of a term, rather than the semantics of a particular expression. Similarly, virtually any term can be used with a deferred reference.<sup>4</sup> This suggests that the phenomenon of deferred reference does not have to do with the semantics of any particular construction. Rather, it involves how we can *use* constructions that have a certain semantics to communicate something different than such constructions semantically express.

A final theoretical consideration that can be brought to bear in arguing that deferred reference is not semantic has to do with the unconstrained nature of any semantic theory adequate to the task. This emerges when one considers the details of the semantic resources one would need to adopt in order to incorporate deferred reference into the semantics. Sag (1981) gives a semantic theory appropriate to the task of incorporating deferred reference into semantic content. Sag introduces “sense-transfer functions” into contexts, and then uses them to interpret expressions

<sup>2</sup> Thanks to Hanna Kim for discussion here.

<sup>3</sup> As Jeff King has pointed out to me, this distinguishes deferred reference from deferred *ostension*. Suppose, pointing at a parked car festooned with tickets, I utter “That driver is going to be upset”. The reference of “that driver” is the driver of the indicated car, even though what I demonstrated is the car. But the driver is still who is denoted by my use of “that driver”, because he satisfies the predicative material “driver”. This distinguishes deferred reference from deferred ostension; the former is not semantic, whereas the latter is.

<sup>4</sup> For example, we can have “Two ham sandwiches are getting irritated”, “Every ham sandwich is clamoring for her check”, “John ham-sandwiched again” (where this latter may mean the same as “John ordered a ham sandwich again”).

in a sentence interpreted relative to a context. On his account, an expression is interpreted relative to a sense-transfer function, which can map the meaning of that expression onto any other meaning. The class of sense-transfer functions is restricted only by pragmatics.

Something like Sag's semantic proposal is required to account for deferred reference. But notice what the resulting "semantic" theory has the power to do. In no sense can it be said that semantic content is "constrained" by conventional meaning. Since, as we just discussed, virtually any word can have a deferred meaning, it follows that any word could in principle acquire any meaning, via a sense-transfer function. The available sense-transfer functions are constrained only by pragmatics. So, the resulting semantic theory is one according to which semantic content is unconstrained by conventional meaning. The semantic content of the word "house" could be the property of being a dog—the only thing that would prevent it from acquiring this semantic content is pragmatic facts about a context.

The moral of this final consideration is that, to capture deferred reference semantically, one would need to adopt a semantic theory where semantic content is not constrained by conventional meaning; in short, an unconstrained semantic theory (that is, constrained only by pragmatics). When capturing a phenomenon within the semantics would result in an unconstrained semantic theory, that suggests that the phenomenon is not semantic. For example, if in order to capture a phenomenon within the semantics, one needs to exploit resources that could allow the semantic content of "Grass is green", relative to a context, be the proposition that snow is white, then the phenomenon is not semantic. This is the principal theoretical reason for denying that deferred reference is semantic.

So, I have given two theoretical reasons for denying that deferred reference is semantic. These considerations are not arguments based on intuitions. As I have already indicated, there is a sense of "intuitive truth-conditions" in which deferred reference enters into intuitive truth-conditions. So one might think that to draw the distinction between semantic content and what is only pragmatic in such a way that the deferred reference of a use of an expression is not part of the semantic content is to abandon the semanticist's view that semantics is the source of our intuitions about the truth-conditions of an utterance.

I don't think that any reasonable way of delineating the border between the semantic and the non-semantic will deliver results that will satisfy all. The responsibility of the semanticist is rather to provide some way of drawing the distinction that preserves the core semanticist claim that the source of our intuitions about truth-conditions is the semantics. Cases like deferred reference are cases in which tutored intuitions *diverge*. It is certainly the case that the non-deferred meaning of (9) is available to all competent users of the language, as in the discourse:

- (18) A: The ham sandwich is getting annoyed.  
 B: That's absurd; sandwiches do not get annoyed.

In such a case, where the putatively literal semantic content is clearly available to all competent users of the language, it is perfectly permissible to let theoretical considerations decide between the putatively literal semantic content and the enriched content (that is, the content enriched with the deferred meaning). This is consistent with the semanticist's position, since this is a case in which speakers have several intuitions easily available to them.

The case of deferred reference contrasts, then, with the case of comparative adjectives. Suppose that a theorist maintained that the semantic contribution of "tall" was something like the semantic content "tall for some comparison class", so that everything in the universe except the smallest thing is tall. Suppose then that I showed a speaker a picture of a tiny dwarfish man, surrounded by normal sized men. Pointing at the dwarfish man, I uttered "That man is tall". On the envisaged theory, the semantic content of my sentence, relative to this context, is a true proposition. The person in question is tall, relative to some comparison class (for example, the class of mice). But this semantic content is utterly inaccessible to the speaker. Unlike the case of deferred reference, there is no possibility of a sensible discourse along the following lines:

- (19) A: (pointing at the dwarfish man) That person is not tall.  
 B: That's absurd; everyone and everything is tall, except for the smallest thing.

What this indicates is that the putative semantic content—that the indicated person is tall for some comparison class—is not available to the competent user of the language. Therefore, it is not consistent with the view I am suggesting to take it as the actual semantic content of the sentence, in context.<sup>5</sup>

I have said that, when the putative semantic content is clearly accessible and tutored intuitions about semantic content diverge, theoretical considerations may enter in to decide where to draw the line between semantic content proper and the rest of what is conveyed in a speech act. As we have seen, deferred reference is one such case. In this case, I gave two theoretical reasons to take the semantic content of a sentence not to be sensitive to deferred meanings. It is instructive to look at another such case in which tutored intuitions may diverge, but theoretical considerations impel us

<sup>5</sup> There are other powerful objections against the view in question. For example, "tall", like other comparative adjectives, is *gradable*. On a degree theoretic view, the function of an intensifier such as "very" is to raise the contextually salient degree of height that something must meet in order to be tall. But, it is mysterious what the semantic function of "very" would be in a sentence such as "Bill is tall, but John is very tall".

to draw a different sort of line between the semantic and the non-semantic; the case of domain restriction.

When a sentence such as (3) or (20) is uttered, we naturally interpret it with respect to a salient domain of quantification:

(20) Every bottle is in the fridge.

For example, (20) could be used to communicate the proposition that every bottle in the house is in the fridge. However, like the case of deferred reference (and unlike the case of comparative adjectives), the unrestricted interpretation is also available to competent language users, as the coherence of the following sort of discourse illustrates:

(21) A: Every bottle is in the fridge.

B: Well, your fridge couldn't possibly be that large! There are bottles somewhere in the world that aren't in your fridge.

So, like the case of deferred reference, though intuitions are sensitive to the domain of quantification, it is nevertheless possible for competent speakers to detect the unrestricted reading of quantified sentences. If, as in the case of deferred reference, there were overwhelming theoretical considerations that mitigated against building the restricted reading of quantified sentences into the semantics, then it would then be acceptable to do so, consistently with the thesis that semantic content delivers intuitive truth-conditions.

However, there are no good theoretical reasons against incorporating domain restriction in the semantics. As we saw, incorporating deferred meaning into semantic content has two disturbing results. First, the semantic content of an expression may be something that does not satisfy the conventional meaning of that expression. Secondly, in order to treat the phenomenon, one needs to employ resources that trivialize the semantics. In contrast, incorporating domain restriction into the semantics brings no such costs.

On the theory of domain restriction advocated in Chapters 2 and 3 of this book, the effect of domain restriction is to restrict the extension of the head noun in a quantified noun phrase. That is, in a sentence such as (19), the effect of domain restriction is just to restrict the interpretation of the property expressed by "bottle", by intersecting its extension with the extension of the property that is the domain restriction. The semantic content of the result will be a subset of the set of bottles. So, the semantic content of the restriction of "bottle" will be something that satisfies the conventional meaning of "bottle". Secondly, in incorporating domain restriction into the semantics, there is no risk of giving the semantics the resources to make "grass is green" express the proposition that snow is white. The only effect context can have is to restrict the interpretation delivered by the conventional meaning of the

head noun in a quantified noun phrase. So, incorporating domain restriction into the semantics is perfectly consistent with the nature of semantic content as intrinsically constrained by conventional meaning.

Since incorporating domain restriction into the semantics does not have theoretical costs, given that domain restriction does affect some level of intuitive truth-conditions, it ought to be incorporated into the semantics. Of course, it is only possible to incorporate domain restriction into the semantics if it is due to semantics, that is, due to the compositional assignment of content to a sentence in context. In previous work, I have argued that there covert structure in quantified noun phrases to which provision of a domain to the semantic content of the sentence containing that noun phrase is due. One argument I have used for this conclusion (in the case of domain restriction as well as other constructions) is what has since been called *the binding argument*. Note that the sentences in (22) are most naturally interpreted as in (23):

- (22) a. In every room, every bottle is in the corner.  
 b. Every student answered every question.
- (23) a. In every room  $r$ , every bottle in  $r$  is in the corner.  
 b. Every student  $x$  answered every question  $y$  on  $x$ 's exam.

One way to generate the readings in (23) is to suppose that there are bound variables in the structure of quantified noun phrases, whose values, relative to a context, generate a domain of quantification.

More specifically, the theory of domain restriction I favor (see Chapters 2 and 3) captures these readings in the following way. Syntactically associated with each nominal are domain restriction indices, of the form “ $f(i)$ ”. Relative to a context, “ $f$ ” is assigned a function from objects to properties, and “ $i$ ” is assigned an object.<sup>6</sup> So, the syntactic structure of the sentences in (23) is similar to the sentences in (24):

- (24) a. Every fireman is tired.  
 b. Every student answered every question.
- (25) a. Every  $\langle \text{fireman}, f(i) \rangle$  is tired.  
 b.  $[\text{Every} \langle \text{student}, f(j) \rangle]$ - $i$  answered every  $\langle \text{question}, f(i) \rangle$ .

<sup>6</sup> To my knowledge, the need for such a function variable in an account of domain restriction was first pointed out in von Stechow (1994, 31). Von Stechow's theory differs from Stanley and Szabó in that his representations associate the domain indices with determiners, rather than nominals.

Given an utterance of, for example, (23b), the speaker intends the value of “f” to be a function from students to their exams, and “j” is bound by the higher quantifier “every student”, yielding the desired reading (23b).<sup>7</sup>

If these are the right representations, then domain restriction is due to the semantics, since it is due to the assignment of values to constituents of a sentence, relative to a context. Evidence that these are the right representations comes from the fact that one detects operator-variable interactions involving quantifier domains, and operator-variable interactions are syntactic in nature.

Of course, binding considerations are certainly not the only way to argue that an allegedly non-semantic phenomenon is due to semantics. For example, the view that the phenomenon in question is non-semantic could be due to an overly simplistic conception of the semantics of some overt expression, and so one way of establishing that the phenomenon is semantic is by giving a more complex semantic clause for some overt expression (cf. again King and Stanley (Chapter 4) on implicature intrusion). Furthermore, binding considerations are not the only way to establish covert structure, since nothing bars the language system from employing syntactic structures containing covert non-bindable indexicals, akin to the overt non-bindable indexicals “I” and “here” of English. But binding considerations are still one way to argue for covert structure, and one that generalizes to a wide range of constructions (see Chapter 1). Because such considerations do provide an argument for the semantic treatment of a wide range of data that pragmaticists have long claimed to be non-semantic in nature, they have recently been widely criticized. In the rest of this chapter, I will look at some of the criticisms of the argument from binding, to see whether they undermine the status of these considerations as arguments for syntactic structure.

### III The Binding Argument

According to the binding argument, if there is a genuine bound reading of a certain construction, that supports the hypothesis that the quantifier in question binds a variable in the syntactic structure of the sentence. For the binding argument to have force, the bound reading must be generated by an expression that is an uncontroversial example of a quantificational expression. The binding argument, *considered as an argument*

<sup>7</sup> The values for the domain indices for the first nominal “student” could be, e.g., the classroom (for “j”) and a function from the classroom to its inhabitants (for “f”). For more discussion of the values of unbound domain indices, cf. Stanley (Chapter 3 of this book).



for syntactic structure, has been interpreted in several different ways. In this section, I discuss the three different ways it has been interpreted.<sup>8</sup>

On the first interpretation of the binding argument, which occurs in unpublished work by Michael Nelson (2001), and in Cappelen and Lepore (2002), the binding argument establishes the existence of covert structure, on pain of ungrammaticality due to vacuous binding. For instance, in the case of (18a), the quantifier “every room” must bind a variable in the syntactic structure of the sentence “every bottle is in the corner”, on pain of ungrammaticality. On the second interpretation of the binding argument, a bound reading of a sentence is evidence for syntactic structure, since bound readings are the semantic effect of a syntactic process (see pp. 48 ff. for details). On this version of the binding argument, it is not potential ungrammaticality that is at issue. Rather, certain kinds of semantic phenomena (for example, bound readings, scope ambiguities) have ultimately a syntactic explanation. On the third (and weakest) interpretation of the binding argument, it is an inference to the best explanation. By postulating a covert variable, one can account for the bound reading, and there is no other satisfactory way to account for it. In Chapter 5, for argument’s sake, I employed this third interpretation in arguing against “free enrichment” accounts of binding.

There are two basic kinds of challenges to the binding argument. First, there are attempts to argue that, whatever the right account of the data, the methodology behind the binding argument is unsound. However, the point of the third version of the binding argument is that merely objecting to the postulation of variables without providing an alternative account is insufficient. It is one thing to raise faults with the methodology, but quite another to provide an account that is equally adequate to the explanatory task. The second kind of response to the binding argument is to attempt to fulfill this obligation, by explaining bound readings without postulating covert structure.

Since the most important task for the person who objects to the binding argument is to explain bound readings without postulating covert structure, I will focus first on accounts that attempt to accomplish this. But before I begin my discussion of such accounts, I want to discuss briefly two approaches to the data that I will *not* discuss at length: variable-free semantics; and free pragmatic enrichment.

<sup>8</sup> Many authors have used bound readings of various constructions to draw disparate morals. Partee (1989) uses bound readings of relational expressions such as “local” and “enemy” to argue that binding is *not* always represented linguistically (thereby drawing the *opposite* conclusion from such data to my conclusion in Chapter 1). Cooper (1993) provides bound readings to argue for the semantic reality of situation variables. Von Stechow (1994) uses bound readings of quantifier phrases to argue that resource domain variables are indexical in nature, but stops just short of arguing that they are syntactically present (*ibid.*, 33). Nevertheless, it is natural to read von Stechow as endorsing that thesis.

A variable-free semantic framework can provide an account of bound readings of sentences without postulating covert structure. There are many different versions of variable-free semantics, but I will briefly focus on the elegant version given in Jacobson (1999). On Jacobson's account, work that might ordinarily be done by postulating syntactic movement or covert structure is done instead by type-shifting in the semantics; a pronoun in the complement of a verb induces a *type-shift* in that verb. A transitive verb has potentially different semantic types, depending upon the number of pronouns that occur within its complement. Complexity in the syntax, on a variable-free account, is replaced by complexity in semantic type assignments to lexical items.

I will not discuss variable-free semantics, because I think the question of whether to implement binding syntactically or semantically is orthogonal to the question at hand, which is whether certain examples demonstrate that intuitive truth-conditions are not generated within the semantics. Both the proponent of variable-free semantics and the more traditional syntactician and semanticist should agree that bound readings of a sentence are of semantic significance. The more traditional syntactician and semanticist should think they are of semantic significance because they indicate hidden syntactic structure, whereas the variable-free semanticist should think they are of semantic significance because they demonstrate that a lexical item is associated with potentially distinct semantic types. Variable-free semantics does not make it easier to argue that certain readings cannot be generated in the semantics; it is irrelevant to this issue.<sup>9</sup>

The second topic I will not discuss is free-enrichment accounts of the data. There are two styles of such accounts. According to the first, when one utters a sentence, via a pragmatic process, the sentence itself is "enriched" into a longer sentence with the addition of lexical material. We may call this *free syntactic enrichment*. According to the second, when one utters a sentence, the semantic content of that sentence (a proposition or a property or propositional function) is enriched by the addition of additional semantic constituents. We may call this *free semantic enrichment*.<sup>10</sup>

<sup>9</sup> An interesting issue arises with what would be captured as free readings of variables in a more traditional framework (Jacobson (1999, 134–5)). On a variable-free framework, there really are no free variables. Explicit pronouns are semantically empty (express the identity function). The effect of a free variable (or a free reading of a relational expression such as "enemy") is to induce type-shifts so that the resulting sentence expresses a propositional function (e.g., in the case of "enemy", a function from persons to singular propositions). On this view, a sentence containing a free variable does not express a proposition, but rather a function from a certain kind of entity (determined by the type of the free variable) to a proposition (or truth-value, depending on one's framework). Satisfying this function is not a matter of free enrichment, but rather closer to what Kent Bach (1994) calls "completion".

<sup>10</sup> Some philosophers of language hold that a sentence expresses a structured semantic content, with specific holes that are saturated by context. I do not consider this to be free semantic enrichment.

I shall not discuss either enrichment account of the intuitive data, not because it is not topical (it clearly is), but because I have discussed such accounts in detail already (Chapter 5). As I have previously indicated, my objection to such accounts is that they over-generate. If free pragmatic enrichment of either kind were a regular mechanism we could appeal to in communication, it would be a mystery why many sentences cannot serve greater communicative functions than they do.

In the next sections, I will rather discuss two challenges to the binding argument that seek to account for the intuitive data without free pragmatic enrichment, either syntactic or semantic. If there were viable alternative accounts of some of the binding data, then that would raise worries about the soundness of the underlying methodology, and thereby threaten to rob the defender of the binding argument of a useful tool by which to establish covert structure in a wide variety of problem cases.

#### IV Binding and Comparative Adjectives

In Chapter 1, the binding argument was used to argue for the syntactic representation of comparison classes for comparative adjectives, such as “tall” and “old”. The target of the arguments there was the following kind of unarticulated constituent clause:

- (R) Relative to any context *c*, “old” expresses the property of being old for a thing of the kind that is salient in context *c*.

In Chapter 1 (p. 55), I pointed out that (R) cannot capture the most natural reading of a sentence such as:

- (26) Most species have members that are old.

The problem with (R) is that it predicts that the occurrence of “old” in (26) must be fixed to a particular species salient in the context of use of (26). But, in the natural reading of (26), the values introduced by the initial quantifier “most species” vary the comparison class to which “old” is applied. (R) cannot account for this reading.

In conversation, I have encountered philosophers challenging this line of argument, by contending that “old” in (26) simply means *old for a thing of its kind*. If so, then (R) produces the desired reading, because the variation is part of the lexical meaning of

In the envisaged process, the role of context is constrained to supply elements of a particular semantic type. Thus, it is conventionally constrained. In contrast, free semantic enrichment is, by its nature, not so constrained. Elements of any semantic type, consistently with the conversational context, could be added.

the adjective “old”. However, this suggestion does not rescue (R), as a similar example shows:

(27) Every sports team has a member who is old.

Intuitively, (27) may express the proposition that every sports team has a member that is old for that sport. But on the view we are considering, “old” expresses the property of being old for  $x$ 's kind; that is  $\lambda x(\text{old for } x\text{'s kind})$ . But each member of a sports team belongs to many different kinds. So it is unclear how to use this suggestion to obtain this reading of (27).

Perhaps we can use this suggestion to emend (R):

(R\*) Relative to any context  $c$ , “old” expresses the property  $\lambda x(x \text{ is old for } x\text{'s } N)$ , where  $N$  is the contextually salient property.

Unlike (R), (R\*) has no trouble with (26). For relative to a context of utterance for (26), the salient property is *species*. So, relative to a context of utterance of (26), “old” expresses  $\lambda x(\text{old for } x\text{'s species})$ , which delivers the correct reading. But (R\*) also promises to help with (27). Relative to a context of utterance of (27), the contextually salient property is *sport*. So, relative to a context of utterance of (27), “old” expresses the property  $\lambda x(\text{old for } x\text{'s sport})$ . According to this clause, then, in a context of utterance of the appropriate sort, (27) expresses the proposition that every sports team has a member who is old for his sport. And this seems to be the desired reading.

However, (R\*) does not work. It faces what we may call *the Bo Jackson problem*. Someone may play more than one sport. In such a case, (R\*) will not deliver a result, since (R\*) requires that there is one unique sport played by each person. Nevertheless, relative to such a situation, (27) may still express a coherent and indeed true proposition, namely the proposition that every sports team  $S$  has a member who is old for the sport played by  $S$ .

Here is a possible repair to (R\*) in light of the Bo Jackson problem:

(R\*\*) Relative to any context  $c$ , “old” expresses the property  $\lambda x(x \text{ is old for some } N \text{ in which } x \text{ participates})$ , where  $N$  is the contextually salient property.

(R\*\*) evades the Bo Jackson problem, since it does not require, of each thing, that it participates in only one kind of the contextually salient property (in the case of (26), only one kind of sport). However, (R\*\*) also fails.

According to (R\*\*), in a context of the appropriate sort, (26) expresses the proposition that every sports team has a member that is old for some sport he plays. Suppose that there are three sports teams, a gymnastics team, a chess club, and a baseball team. One person, Bob, plays for all three teams. Bob is old for a gymnast but not old for a chess player or a baseball player. No other members of the teams are old for their sports. Intuitively, what an utterance of (27) expresses, relative to

this situation, is false. However, according to (R\*\*), the proposition expressed by (27) should be true in this situation. For each sports team does have a member who is old for some sport he plays. Each sports team contains Bob, who is old for a gymnast.

So there does not seem to be any easy repair of a rule such as (R). If one wishes to capture semantically all of the intuitive judgements about truth and falsity we have discussed, examples such as (26) and (27) seem to require a syntactically represented comparison class (or some other mechanism that imitates binding).

## V Quantifying over Contexts

One classic example of an unarticulated constituent analysis of a construction involves example (28), from Perry (1986):

(28) It's raining.

(U) "It is raining(*t*)" is true in a context *c* if and only if the denotation of "rain" takes  $\langle t, l \rangle$  to the true, where *l* is the contextually salient location in *c*.

In Chapter 1 (pp. 52 ff.), I used binding considerations against an unarticulated constituent analysis like this. In particular, I used examples such as:

(29) Every time John lights a cigarette, it rains.

The unarticulated constituent analysis suggested in (U) cannot derive the natural reading of (29), where the location of the raining varies with the values of the initial quantifier, "every time John lights a cigarette".

My purpose in giving this argument was not so much to advance my own account of such examples as to reject an unarticulated constituent analysis. But I did propose two positive "articulated" accounts of the data (p. 53). According to the first, "rain" is associated with an event or situation variable, which is bound by the initial quantifier "every time". According to the second, "rain" is associated with a pair of variable positions, one of which determines a time, and the other a location, both bound into by the initial quantifier.

Peter Pagin (2005) seeks to evade the need for either kind of analysis, by treating quantifications such as "every time John lights a cigarette" as quantifiers over contexts. On Pagin's analysis, (29) ends up having the truth-conditions in (30):

(30) For every context *c'* differing from *c* at most in its time and location indices, "if John lights a cigarette, then it rains" is true in *c'*.

On Pagin's treatment, there is no need for a variable for events or locations, because the initial quantifier is over contexts.

My concern with Pagin's analysis is a familiar one with operators that shift contextual features, noted originally by Lewis (1981, 86):

we need to know what happens to the truth values of constituent sentences when one feature of context is shifted and the rest are held fixed. But features of context do not vary independently. No two contexts differ by only one feature. Shift one feature only, and the result of the shift is not a context at all.

Suppose I utter (29) in a context *c*. I am the speaker of *c*. But to obtain the right truth-conditions, we need to quantify over all *n*-tuples that differ from *c* only in their location and time indices. But some *n*-tuples of indices will not be possible contexts of use.

To make Lewis's point vivid, consider the sentences in (31)

- (31) a. Whenever I'm politely listening to someone speaking, it starts to rain.  
 b. Whenever wind blows through a mountain pass, it starts to rain.

Pagin's truth-conditions for the sentences in (31), considered as uttered in a context *c*, are:

- (32) a. For every context *c'* differing from *c* in at most its location and time indices,  
 "If I'm politely listening to someone speaking, it rains" is true in *c'*.  
 b. For every context *c'* differing from *c* in at most its location and time indices,  
 "If the wind blows through a mountain pass, it starts to rain" is true in *c'*.

Now consider the contexts *c'* involved in obtaining the correct truth-conditions. To obtain the right truth-conditions, many of these must be packages of indices that are not possible contexts of use. For example, to obtain the right truth-conditions for (31a) via (32a), we need contexts in which the speaker in *c* is the addressee in *c'*, rather than the speaker. But these will not be contexts that differ from *c* only in their location and time indices. A different problem surfaces for (31b). A context is one in which the agent of the context is at the time and location of the context. So Pagin predicts that (31b) is true just in case whenever the speaker in *c* is at the time and location of the relevant mountain pass, it starts to rain when the wind blows through. But clearly, these truth-conditions are too weak. (31b) would be falsified if there are situations with no one around (and hence no agents) in which the wind blows through a mountain pass and it doesn't start to rain at that location.<sup>11</sup>

So, I am skeptical that appealing to quantifiers over contexts will help in accounting for bound readings of alleged unarticulated constituents. The problem is that

<sup>11</sup> Pagin discusses a similar problem (see the discussion surrounding principle (I)). But his discussion cannot accommodate (31b), since his approach involves quantifying only over contexts, and to obtain the right truth-conditions for (31b), one needs to quantify over indices that are not possible contexts of use.

quantifying over contexts results in truth-conditions that are too weak, given the paucity of contexts of use.

## VI The Challenge from Over-Generation

In the past two sections, I have discussed attempts to capture the binding data without postulating unpronounced structure. I now turn to challenges to the binding argument. The most common sort of objection involves *over-generation*.<sup>12</sup> According to this kind of challenge, if one postulates variables when bound readings are available, what results is an over-generation of variables in the syntax.

In what follows, I will respond to the over-generation concern for the binding argument. But first, I want to note an oddity about the strategy of advancing over-generation objection against the binding arguments. Those who advance such objections typically do so in support of *pragmatic* accounts of the bound readings. But pragmatic accounts of the data, as I have emphasized, over-generate more than any other account possibly could. For pragmatic accounts are, by their nature, *unconstrained*

<sup>12</sup> One different type of objection to the binding argument, discussed in Cappelen and Lepore (2002, 276–7) is that the variables the binding argument would have us postulate behave differently than overt pronouns, in particular in anaphora. For example, as they point out, it is odd to follow up an utterance of “Tigers are mammals”, by “and it is a big domain” (with the “it” referring to the domain associated with “tigers”). This objection deserves more attention than I can give it here (unpublished work by Adam Sennet and Brett Sherman is important in this regard). But one reason to be suspicious of the argument is that it would apply to an alarmingly large range of constructions. Of course, Cappelen and Lepore revel in this fact; as they note (*ibid.*, 279), these considerations would also tell against postulating variables for comparison classes for comparative adjectives (cf. also Cappelen and Lepore (2005)). Their argument would also tell against postulating variables for degrees for adjectival constructions, as witnessed by the oddity of “John is tall, and it is a high degree”. More problematically, the considerations also entail that the implicit anaphoric elements associated with relational expressions such as “local” and “enemy” are not syntactically realized. For example, suppose Bill utters “John talked to an enemy in 2004”, thereby expressing the proposition that John talked to an enemy of Bill in 2004. It is not possible for someone to follow this utterance up by saying “He has many enemies”, where the “he” is genuinely *discourse anaphoric* on the covert variable that refers to Bill. Similarly, suppose Bill utters “John talked to an enemy”, meaning an enemy of the United States. It is not possible to follow this up with, “And it is a big country”, where “it” is *discourse anaphoric* on the covert variable. So, if this argument were correct, *implicit anaphora* would not be syntactically realized. Similarly, it is plausible to take epistemic modals to involve implicit anaphora; the occurrence of “might” in a token of “It might be raining in Paris on 19 July 2004” is to be taken relative to the knowledge state of the person making the utterance. But one cannot follow up someone’s utterance of “It might be raining in Paris” with “He is strange”, where “he” is an anaphoric pronoun (contrast this with the acceptability of following “According to John, it might be raining in Paris” with “He is strange”, where “he” is uttered with anaphoric intent). The argument therefore proves too much, unless Cappelen and Lepore are also willing to use it to reject the syntactic representation of implicit anaphora.

by linguistic meaning. Were some pragmatic account to be correct, there would be numerous sentences that would allow readings that they actually do not allow (Chapter 5).<sup>13</sup>

The first over-generation objection I will discuss is due to Cappelen and Lepore (2002, 273):

A confused mathematical anthropologist (call her “Sally”) trying to find out if mathematical truths are universal utters (5) as a summary of her findings:

(5) Everywhere I go,  $2 + 2 = 4$

Here’s the binding argument applied to (5):

Intuitively, (5) says that for every place Sally goes,  $2 + 2 = 4$  at that place. So we should present the logical form of (5) along the following lines:

(5\*) For all places  $x$ , if Sally goes to  $x$ , then  $2 + 2 = 4$  at  $x$ .

The quantifier phrase “Everywhere Sally goes” is binding a place variable in the logical form of “ $2 + 2 = 4$ ”—otherwise, there would be nothing for the quantifier phrase to bind. This establishes that the logical form of the sentence “ $2 + 2 = 4$ ” has a freely occurring place variable.

Since there is obviously *no* variable ranging over locations in “ $2 + 2 = 4$ ”, this is a *reductio* of the binding argument.

Before responding to this argument, we should be clear about what version of the binding argument Cappelen and Lepore have in mind. As I discussed above, there are three versions of the binding argument. According to the first, it has to do with grammaticality: one must postulate a place variable in the logical form of a sentence, or else what cannot explain the grammaticality of the larger construction. According to the second: it does not have to do with grammaticality, bound readings are taken to be a reflection of syntactic binding. According to the third reading: it is an inference to the best explanation of the bound reading.

Cappelen and Lepore, in their paper, address the first of these versions of the binding argument. Since I am not aware of that version being promoted in published

<sup>13</sup> Indeed, in the thousands of pages that have been written over the last decade arguing for pragmatic (non-semantic) accounts of a wide range of apparently semantic phenomena, I am not aware of a single attempt to provide a response to the threat of over-generation to pragmatic theories. Indeed, I am not even aware, aside from passing footnote references, to a discussion of the over-generation threat facing such theories. Given this silence, there is some irony involved in such theorists’ extreme sensitivity to over-generation worries with alternative positions. I hope that the sensitivity such theorists evince to over-generation objections will soon be reflected in greater attention to these worries with their favored accounts.



work, I am not sure it should be the focus of critical attention.<sup>14</sup> That is, the obvious *grammaticality* of example (5) poses no worries for the advocate of the binding argument. So I take it that the feature of the example that is supposed to concern the advocate of the binding argument is the claim that the intuitive reading of the example is (5\*). If the intuitive reading of (5) is (5\*), then it would seem that the advocate of the binding argument is committed to postulating a place variable in the logical form of “ $2 + 2 = 4$ ”.

However, I do not see that (5\*) is the intuitive reading of (5), and I do not see that Cappelen and Lepore even believe that (5\*) is a reading of (5). As Cappelen and Lepore (*ibid.*) point out, “it is close to indisputable that arithmetical statements lack hidden indexicals referring to places”. Presumably, the reason they are so convinced of this is that it is unclear what it even means to speak of an arithmetical statement being true *at a place*. If that is the intuition, then it is equally hard to see how (5\*) is a legitimate reading of (5). Sally may intend (5\*) as a reading of (5), because, as Cappelen and Lepore put it, she is a “confused mathematical anthropologist”. But the fact that someone confusedly believes a sentence has a certain reading does not give that sentence that reading.

Of course, it is uncontroversial that (5) is grammatical. But nothing follows from this, other than a rejection of the first version of the binding argument. But since the first version of the binding argument is not one that has ever been advanced in print, it is not germane to the issue.

Brehehy (2004) has leveled another sort of over-generation objection against the binding argument, this involving what he calls the problem of multiple dependencies. Since Brehehy’s arguments are interesting and illustrate important points, it is worth going over them in detail.

Recall that on my favored account of domain restriction, motivated by binding considerations, each noun is syntactically associated with two indices, a function index and an argument index. Relative to a context, the function index is assigned a function from objects to properties, and the argument index an object. So, if I have New Jersey in mind when I say “Every politician is saintly”, then New Jersey is the value of the object index, and perhaps the function index is assigned a function from states to the property of being an inhabitant of that state. It is this view that is Brehehy’s target.<sup>15</sup>

<sup>14</sup> Cappelen and Lepore cite Nelson (2001), who seems to have the first interpretation of the binding argument in mind. But the relevant passages in Nelson (2001) involve a summary and subsequent critique of Chapter 1, and in that paper, I certainly did not have the first interpretation of the binding argument in mind.

<sup>15</sup> In previous work (Chapters 2 and 3) I took the function index and the object index to “co-habit” a node with the head noun. I no longer think this is correct. Talk of “co-habiting a node” with a lexical item suggests that domain indices are part of lexical structure. But the position that domain indices are merely part of lexical structure is not consistent with the general

The first sort of example Breheny discusses is:

- (33) Every student was feeling particularly lucky and thought no examiner would notice every mistake.

Breheny argues that (33) has the reading in:

- (34) [Every student]<sub>x</sub> thought [no examiner]<sub>y</sub> would notice [every mistake made on a paper x turned in and y examines]<sub>z</sub>.

Breheny argues that this sort of example is problematic, because the approach I advocate cannot generate reading (34) of (33), without postulating more syntactic variables. That is, Breheny worries that the methodology of the binding argument commits its proponent to postulating implausibly many variables, certainly more than just the domain variables discussed above.

But I do not see the worry with this particular example. The noun “examiner” is a relational noun. It is associated (cf. Stanley (2000)) with a syntactically represented index. This fact straightforwardly generates the desired reading. That is, an independently motivated claim about relational expressions, together with an independently motivated theory of domain restriction, straightforwardly predicts that (33) has the reading:

- (35) [Every student]<sub>x</sub> thought [no examiner, x]<sub>y</sub> would notice [every mistake  $f(x)$ ]<sub>z</sub>.

We may assume “ $f$ ” is assigned a function from students to their exam questions. So, we can straightforwardly predict a reading of (33) according to which every student thought no examiner of that student would notice every mistake on that student’s exam. So recognizing the relational nature of “examiner” enables us to derive the natural reading with only independently motivated resources.

Breheny also has examples that do not involve relational nouns. However, they also do not raise problems for any view I have defended. Consider Breheny’s “multiple dependence” example (36a), which he claims to have reading (36b):

- (36) a. Every paranoid artist thinks no dealer will stop at selling every forged painting.  
 b. [Every paranoid artist]<sub>x</sub> thinks [no dealer]<sub>y</sub> will stop at selling every [[forged [painting by x]] coming into y’s possession.

view (advocated in Chapter 1) that considerations such as binding, weak-crossover, and ellipsis are evidence for genuine syntactic structure, rather than mere lexical structure. Secondly, both Breheny (ibid.) and Timothy Williamson (2004) provide evidence that the domain index must sometimes be outside the scope of adjectives modifying the head noun. Since the domain indices must sometimes be inside the scope of modifying adjectives (cf. Stanley (Chapter 3 of this book)), this suggests that the domain indices occupy their own terminal nodes that can have different adjunction sites.

On my account, every noun (or  $N'$ ) in every Quantified Noun Phrase (QNP) is associated with a domain index. So, (36a) is predicted to have a reading as in (37):

- (37) [Every paranoid artist]<sub>x</sub> thinks [no dealer  $f(x)$ ]<sub>y</sub> will stop at selling [every forged painting  $f(y)$ ].

Relative to the envisaged context, “ $f$ ” is assigned a function from artists to people in possession of forged paintings of that artist, and “ $f$ ” is assigned a function from dealers to their collections. So, the theory predicts that (36a) has the reading:

- (38) Every paranoid artist thinks that no dealer in possession of forged paintings of that artist will stop at selling every forged painting in their collection.

Intuitions are subtle here. But both of my informants think that (36a) clearly has reading (38), and none of my informants thinks that (36a) has reading (36b). In other words, there is a strong difference among my informants between (36a) and (39):

- (36) a. Every paranoid artist thinks no dealer will stop at selling every forged painting.

- (39) [Every paranoid artist]<sub>x</sub> thinks no dealer will stop at selling every forged painting of his<sub>x</sub>.

My informants do not obtain reading (39) of (36a). They only obtain reading (38) of (36a). And that is precisely what an account that only postulates a single pair of domain variables (one argument and one function variable) would suggest.

A final example of Breheny is:

- (40) Every company knows that none of the pension fund can be diverted away from any former employee.

Breheny claims this example has the following reading:

- (41) [Every company]<sub>x</sub> knows that [none of the pension fund]<sub>y</sub> can be diverted away from any [[former employee of  $x$ ] who is due some of  $y$ ].

But again, no additional variables are needed to capture the appropriate reading semantically. Since “employee” is a relational noun, (40) is predicted to have the reading in (42):

- (42) [Every company]<sub>x</sub> knows that [none of the pension fund  $f(x)$ ]<sub>y</sub> can be diverted away from any former [employee,  $x$ ].

Relative to the envisaged context, “ $f$ ” is assigned a function from (for example) companies to their benefit programs. So (41) expresses the proposition that every company knows that none of the pension funds of that company can be diverted away from any former employee of that company. And this is precisely the desired reading.

Brehehy writes that “if we wanted to pursue a variable-rich approach, given these multiple dependencies, we seem to need to assume that QNP structures contain a plethora of hidden variables at different levels which are vacuously assigned (to what?) when not used.” But, as I have argued, this worry is unfounded. The assumption that relational nouns are associated with syntactically realized implicit arguments, together with the assumption that QNPs are each associated with a single domain index, is sufficient to explain all the data. Indeed, the framework even explains why certain sentences *lack* readings, as in reading (36b) of (36a).

There is a final sort of over-generation objection against the binding argument that I wish to discuss. Because I think this sort of objection hinges on a confusion about the methodology of semantics, responding to it involves less focus on empirical detail, and more on foundational matters.

According to Francois Recanati, the binding argument involves what he calls the “binding criterion”. The binding criterion is that intuitively bound readings must be reflected by bound variables in the syntax. Given his assumption, here is his objection (Recanati (2004), 106–7):

We can say:

- (15) John is anorexic, but whenever his father cooks mushrooms, he eats.

On a natural interpretation, we understand that John eats the mushrooms his father has cooked. Intuitively, a form of binding is operative here; for the food eaten by John co-varies with the food cooked by his father. Such examples show that intuitive binding, *per se*, does not entail the existence of a free variable in logical form. The Binding Criterion, on which Stanley’s argument rests, must be rejected.

Recanati’s argument against the binding criterion has two premises. The first premise is that the intuitive reading of his sentence (15) is that whenever John’s father cooks mushrooms, John eats the mushrooms his father cooks. Recanati’s second premise is that there is no covert variable for *what John eats* in the logical form of “he eats”. Assuming these two premises, Recanati concludes that the binding criterion, which requires such a covert variable, must be rejected.

Recanati’s example is deliberately modeled upon examples discussed in Chapter 1, such as:

- (43) Whenever John lights a cigarette, it rains.

There, I concluded that “it rains” does contain a variable that can be bound by the initial quantifier, “whenever John lights a cigarette”. The evidence that it can be is due to the fact that the intuitive reading of (43) is:

- (44) Whenever John lights a cigarette, it rains at the location of that lighting.

I also suggested various ways of accounting for the bound reading in question (p. 53). Recanati's suggestion is that his example (15) is analogous to (43), so that whatever methodology leads one to postulate an unpronounced variable in "it rains" to account for the bound reading in (43) should lead one, incorrectly, to postulate an unpronounced variable for what is eaten in "John eats".

However, Recanati's (15) is simply not analogous to (43). In particular, the first premise of Recanati's argument is false (whereas the corresponding premise for (43) is true). As Luisa Marti (ms) has shown, it is not the case that the intuitive truth-conditions of (15) are what Recanati says they are. The following two discourses, due essentially to Marti, reveal the clear difference between (15) and (43):

- (45) A. Whenever John's father cooks mushrooms, John eats.  
 B. #No he doesn't—he eats broccoli when his father cooks mushrooms.
- (46) A. Whenever John lights a cigarette, it rains.  
 B. No it doesn't—though it rains somewhere else.

As the oddity of Marti's discourse in (45) clearly demonstrates, it is not permissible to deny the content of A's assertion on the grounds that one thinks that John eats broccoli, rather than mushrooms, when his father cooks mushrooms. This demonstrates that the intuitive truth-conditions of an utterance of Recanati's (15) is not that whenever John's father cooks mushrooms, John eats the mushrooms his father cooks. It is rather that whenever John's father cooks mushrooms, John eats *something*. In contrast, the acceptability of the discourse in (46) demonstrates that the intuitive content of (43) is (44).

According to the binding argument, when a bound reading is part of the intuitive truth-conditions of an utterance, it is the result of a quantifier–variable interaction. The problem with Recanati's example is that (as Marti's argument demonstrates), it is no part of the intuitive truth-conditions of an utterance of (15) that John ate the mushrooms his father cooked. Our intuitions about truth and falsity clearly reveal this to be a reflection of our background assumptions, combined with the semantic content of (15), which is just that in whatever situation John's father cooked mushrooms, John ate something in that situation. In contrast, our intuitions about the truth-conditions of utterances of sentences such as (43) clearly reveal that relativity to a location parameter is part of their intuitive truth-conditions.

The distinction between a verb like "rains" and a verb like "eats" can be seen even in non-embedded sentences. Suppose Bill has cooked a mushroom dinner. Pointing at a dirty plate on the table, and intending to communicate that John has eaten the mushrooms Bill ate, I utter:

(47) John ate.

Suppose one knew that John had just eaten, but he did not eat the mushrooms Bill cooked. It is still clearly not permissible to follow my assertion with:

(48) No he didn't; he ate broccoli instead.

In contrast, suppose that it is raining in New York City, where I am located. Speaking on the phone with Delia, who is in Ithaca, I utter:

(49) It's raining.

Delia, seeing on television that the sky is clear in New York, utters:

(50) No it isn't. But it's raining here.

Delia's reply is perfectly acceptable, in contrast to the unacceptable (48). This demonstrates that the location is part of the intuitive truth-conditions of an utterance of (49), whereas what is eaten, even when it is salient, is clearly not part of the intuitive truth-conditions of an utterance of (47).

Of course, when we hear an utterance of Recanati's (15), we are liable to assume that John ate the mushrooms his father cooked. The reason we assume this has nothing whatever to do with semantics. Rather, we assume that when your father is cooking a meal, and you are eating together with him, it is expected behavior to eat what your father has cooked. We also assume, when we hear an utterance of Recanati's (15), that John's father is not a Martian. But it would be absurd to suppose that it is the semantics that "tells" us that John's father is not a Martian. It is similarly absurd to suppose that information we acquire via the background assumption that people generally eat the meals that are cooked for them must be supplied semantically, if the semanticist is right.<sup>16</sup>

Similar mistakes to these have been made by other pragmaticists, in their discussions of the view of Chapter 1 that the source of the intuitive truth-conditions of an utterance are covert structure. For example, Wilson and Sperber (2002) exploit the example:

(51) I must wash my hands: I've eaten [using my hands, rather than, say, being spoon-fed].

<sup>16</sup> Iliria Frana (ms) provides an argument against the binding argument that is also undermined by these considerations. Frana considers the sentence "Paolo is a real curious guy; every time he finds something, he opens it". According to Frana, an utterance of this sentence has the intuitive truth-conditions that Paolo always opens the thing he finds in a manner appropriate to that thing. But intuitively what is said would not be *false* if Paolo always opened what he found in a manner that was not appropriate to that thing. Therefore, this is no part of the intuitive truth-conditions of Frana's sentence. A similar point dispenses with the reply to the binding argument in Stalnaker (2004, 110–11).

According to Wilson and Sperber, the bracketed material is part of the intuitive truth-conditions of an utterance of “I must wash my hands: I’ve eaten”. So they conclude that anyone who defends the clear and elegant explanation of the source of intuitive truth-conditions is committed to the view that “eat” has an argument place for *the manner of eating* (see also Carston (2002, 203–4)). But it is false that the manner of eating, even in the example they envisage, enters into the intuitive truth-conditions, as the oddity of Bill’s utterance in the Marti-style discourse in (52) demonstrates:

- (52) John: I must wash my hands: I’ve eaten.  
 Bill: No you didn’t; you got spoon-fed.

Of course, we would naturally assume, given John’s discourse, that he ate with his hands. We would also assume that he was not from Mars, and that he was not the product of in-vitro fertilization. It is just as absurd to take it as an objection to the semanticist that such information is not provided by the semantics of the verb “eat” as it is to take it as an objection to the semanticist that the semantics does not provide for a manner of eating. Such information has nothing to do with intuitive truth-conditions of an utterance.

The responsibility of the semanticist is rather to show that speaker intuitions about the truth-conditions of an utterance are due to semantics. Unfortunately, too many objections to the semanticist assume that the responsibility of the semanticist is to generate within the semantics all information that a competent speaker and member of a culture may derive from a communicative act. Such objections seem to presume, absurdly, that the semanticist’s position is incompatible with Grice. As Marti’s tests show, being more subtle about judgements of truth and falsity can clearly reveal the distinction between what is part of intuitive truth-conditions proper and what is conveyed by the communicative act to a hearer who combines these truth-conditions with her background knowledge about the world.

## Conclusion

My purpose in this chapter has been to defend the claim that the intuitive truth-conditions of an utterance are due to semantic interpretation. Many of those who have objected to it have done so by saddling the position with absurd theoretical commitments, such as the position that all information conveyed in any discourse is due to the semantics. Part of my goal in this chapter has been to explain what costs the semanticist’s position incurs, and what costs it does not, by elucidating the target concept of semantics. As I have argued, in certain cases (for example, that of deferred reference), the semanticist must make decisions about the defendum that

are informed by theoretical considerations. But this is the ordinary practice even in the human sciences.

## References

- BACH, KENT (1994), 'Conversational Implicature', *Mind and Language*, 9: 124–62.
- BREHENY, RICHARD (2004), 'A Lexical Account of Implicit (Bound) Contextual Dependence', in *Semantics and Linguistic Theory (SALT)*, 13.
- CAMP, ELIZABETH (ms) 'Lodging in the Words'.
- CAPPELEN, HERMAN, and LEPORE, ERNIE (2002), 'Indexicality, Binding, Anaphora, and A Priori Truth', *Analysis*, 62/4: 271–81.
- CAPPELEN, HERMAN, and LEPORE, ERNIE (2005), 'A Tall Tale in Defense of Semantic Minimalism and Speech Act Pluralism', in Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 197–219.
- CARSTON, ROBYN (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication* (Oxford: Blackwell).
- COOPER, ROBIN (1993), 'Generalized Quantifiers and Resource Situations', Aczel, Israel, Katagiri, and Peters (eds.), *Situation Theory and its Applications* (Stanford: Stanford University Press).
- FRANA, ILARIA (ms), 'The Constituents of Meaning'.
- JACOBSON, PAULINE (1999), 'Towards a Variable-free Semantics', *Linguistics and Philosophy*, 22: 117–84.
- LEVINSON, STEPHEN (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (Cambridge: MIT Press).
- LEWIS, DAVID (1981), 'Index, Context, and Content', in Stig Kanger and Sven Ohman (eds.), *Philosophy and Grammar: Papers on the Occasion of the Quincentennial of Uppsala University* (Dordrecht: Reidel), 79–100.
- MARTI, LUISA (ms), 'Unarticulated Constituents Revisited'.
- NELSON, MICHAEL (2001), 'When is an Expression Context-sensitive?', unpublished ms.
- PAGIN, PETER (2005), 'Compositionality and Context', in Gerhard Preyer and Georg Peter, *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 303–48.
- PARTEE, BARBARA (1973): 'Some Structural Analogies Between Tenses and Pronouns in English', *Journal of Philosophy*, 70/18: 501–9.
- PARTEE, BARBARA (1989), 'Binding Implicit Variables in Quantified Contexts', *CLS 25: Papers from the 25th Annual Regional Meeting of the Chicago Linguistics Society* (Chicago: University of Chicago Press), 342–65.
- PERRY, JOHN (1986), 'Thought without Representation', in Supplementary Proceedings of the Aristotelean Society, 60: 137–52.
- RECANATI, FRANÇOIS (2004), *Literal Meaning* (Cambridge: Cambridge University Press).
- SAG, IVAN (1981), 'Formal Semantics and Extra-linguistic Context', in Peter Cole (ed.), *Syntax and Semantics*, 13: Radical Pragmatics (New York, Academic Press), 273–94.



- STALNAKER, ROBERT (2004), 'Comments on "From Contextualism to Contrastivism"', *Philosophical Studies*, 119/1 and 2: 105–17.
- STERN, JOSEF (2000), *Metaphor in Context* (Cambridge: MIT Press).
- VON FINTEL, KAI (1994), 'Restrictions on Quantifier Domains', diss., University of Massachusetts.
- WILLIAMSON, TIMOTHY (2004), 'Everything', in *Philosophical Perspectives*, 26: *Philosophical Linguistics and Philosophy of Language*.
- WILSON, DEIRDRE, AND SPERBER, DAN (2002), 'Truthfulness and Relevance', *Mind*, 111: 583–632.

## Review of François Recanati's *Literal Meaning*

In philosophy, the twentieth century began with the thought that the context-dependence and vagueness of natural language undermined the possibility of providing a systematic account of the meaning of natural language sentences. Philosophical reflection on language continued through its middle period with an even more explicit emphasis on the unsystematic character of language. But then Paul Grice showed how to explain some of the unruly effects of context on linguistic communication by appeal to general conversational principles. From the other direction, Richard Montague and his students showed that much of what appeared to be unsystematic was in fact explicable. Indeed, Montague and his coterie approached this task using the tools developed by the descendants of the very philosophers who had despaired of the possibility of providing a rigorous semantics for natural language. As syntax and semantics became increasingly sophisticated, vagueness and context-dependence became objects of formal study, rather than phenomena whose existence demonstrated the impossibility of such work. From the perspective of those fluent with the tools of Chomsky, Grice, and Montague, conclusions from premises about the unsystematic nature of natural language began to look a bit like a previous era's skepticism about the possibility of a systematic physical theory of the universe.

It is somewhat surprising, then, that philosophy of language at the end of the twentieth century, and the beginning of the twenty-first, has been dominated by a wealth of papers and books seeking to return us to the pessimistic conclusions of the past. François Recanati is one of the major figures in this literature. In Recanati (1993, 227–74), he argued that what is intuitively said by an utterance is affected by context in ways that could not be explained by any combination of Chomsky, Montague, and Grice (that is, ordinary syntax and semantics, together with Gricean pragmatics). Since the publication of that work, he has been developing this thesis in detail. His arguments for the thesis he calls *contextualism* are brought together in characteristically clear and concise form in *Literal Meaning*. According to contextualism (4), “the contrast between what a speaker means and what she says is illusory, and the notion

of ‘what the sentence says’ is incoherent”. *Literal Meaning* is devoted to defending the thesis of contextualism against rival views, and exploring some of its consequences for particular linguistic phenomena.

For Grice, implicating a proposition was a species of conscious intentional action. A speaker uses a sentence, and thereby intentionally expresses one proposition in order to implicate another. For Grice, then, speakers are aware of the propositions they express when they utter sentences. But according to Recanati’s Contextualism, virtually any proposition we are aware of asserting is one that is thoroughly affected by what he calls “primary pragmatic processes”. So Recanati holds that Grice’s notion of what is said by an utterance (the input to implicature) is a thoroughly non-semantic notion.

There are a number of other theorists who agree with Recanati that Grice’s notion of what is said by an utterance cannot be what is delivered by a systematic semantic theory. First, there are other contextualists. Relevance theorists, such as Robyn Carston, Dan Sperber, and Deirdre Wilson, agree with Recanati that any notion of what is said is thoroughly pragmatic in nature. Relevance theorists also believe that there is a systematic account of the sort of pragmatic inferences by means of which interpreters grasp what is said in context. Other contextualists, such as John Searle and Charles Travis, are grim pessimists who reject the possibility of any such systematic account. Then, there is an apparently non-contextualist position, which Recanati calls the *Syncretic View*. Adherents of the Syncretic View agree with contextualists that what speakers consciously intend to express by their utterances is not usually the semantic content of the sentences uttered (even relative to that context of use). But adherents of the Syncretic View nevertheless maintain that sentences, relative to a context of use, do have a semantic content, and that the fact that they do is in some sense important for the theory of meaning. Both contextualists and adherents of the Syncretic View agree, however, that what is consciously available to the speaker as what she primarily intended to express by her utterance is not in general what is delivered by the semantic interpretation of the sentence she uttered (even relative to that context of use).

The Syncretic View is a position particularly popular among philosophers of language, in part for sociological reasons. Philosophy of language in the 1980s was dominated by disputes between Millians, who held that the semantic content of a name was exhausted by the object in the world to which it referred, and those who maintained that there is some other semantic content associated with a name, for example a mode of presentation of the object to which that name referred. Millians in the 1980s defended the thesis that the semantic content of the sentence “Hesperus is Phosphorus” (relative to a context of use) expressed the same proposition as the semantic content of the sentence “Hesperus is Hesperus”, despite our inclination to believe otherwise. Similarly, Millians believed that “John believes that Hesperus is

Phosphorus" expressed the same proposition, relative to a context of use, as "John believes that Hesperus is Hesperus", despite the fact that ordinary speakers who assertively utter the first generally believe that they are saying something different than they would be saying if they had uttered the second. So Millians are antecedently committed to the Syncretic View, since they think that sentences do express semantic contents, and that the semantic content a sentence has relative to a context can be quite distinct from the content the speaker intends to assert (or thinks she asserts) by her utterance of that sentence in that context. Those who upheld the Millian line in the 1980s have spent some portion of the ensuing period defending their Syncretic commitments (see the work of Nathan Salmon and Scott Soames). But Millians and their descendants are not the only defenders of the Syncretic View among philosophers. Kent Bach has argued for years for a certain version of the Syncretic View, where the semantic content of many sentences, even relative to a context of use, is not a full proposition. Most recently, Herman Cappelen and Ernie Lepore have argued for a version of the Syncretic View, where semantic contents of sentences are always propositional.

Recanati devotes the third chapter of his book to comparing and contrasting his view with relevance theory, and the fourth chapter criticizing the Syncretic View; I will discuss only the latter chapter here. Advocates of the Syncretic View hold that the proposition intuitively expressed by an utterance is only rarely the semantic content of the sentence relative to the context of utterance (King and Stanley (Chapter 4) call the Syncretic View *semantic modesty*). The central problem for the Syncretic View is that the notion of semantic content appealed to in the theory threatens to be an idle wheel in an explanation of linguistic practice. King and Stanley (pp. 145 ff.) argue, as against those who think that complex expressions have *characters* in the sense of David Kaplan (roughly standing linguistic meanings), that the characters of complex expressions play no explanatory role in an account of meaning that cannot be played simply by appealing to the characters of lexical items and the syntactic structure of sentences. Borrowing this line of thought, Recanati argues (64) that what work is done by postulating the minimal proposition expressed by a sentence in a context can be done simply by appeal to the contents of individual words relative to that context, and the syntactic structure of the sentence.

I am no fan of the Syncretic View. Furthermore, as will become clear below, I suspect that Syncretism is not so different from Recanati's favored version of the thesis he calls "contextualism". Nevertheless, I suspect that Recanati cannot straightforwardly borrow King and Stanley's criticism of the view that complex expressions have characters as an objection to all forms of the Syncretic view. For King and Stanley's point is that (for example) sentence-level characters play no role in a semantic theory. In particular, King and Stanley assert that there are no operators that take sentence-level characters as objects. In contrast, some advocates of the Syncretic View

presumably *do* think that there are operators that take minimal propositions as objects. I suspect that Nathan Salmon and Scott Soames (or at least relevant time-slices of them) think that modal operators such as “it is necessary that” and “it is possible that” take minimal propositions as objects. This suggests that the minimal proposition does some semantic work on some versions of the Syncretic View.<sup>1</sup>

There are two versions of the Syncretic View. According to the first version, which I will call *propositional syncretism*, semantic contents of sentences in contexts are always propositions, but not usually the propositions the users of those sentences intend primarily to assert. Rather, semantic contents are generally “minimal propositions” (as in the work of Cappelen and Lepore (2005)). According to the second version, which I will call *non-propositional syncretism*, semantic contents of sentences in context are occasionally non-propositional (as in Bach’s work). Recanati provides some persuasive criticisms of propositional syncretism. For example, according to what Recanati calls the “common denominator” approach to the minimal proposition, the semantic content of a sentence *S* in context *c* is what would be asserted and conveyed in *every* normal context *c'* in which the reference of all indexicals in *s* is the same as their reference in *c* (see Soames (2002, 106)).<sup>2</sup> The common denominator approach is behind all recent presentations of propositional syncretism (see *ibid.* (56–63), Cappelen and Lepore (2005, 57, 143)). The idea behind the common denominator approach is that the semantic content of a sentence relative to a context is the minimal propositional information that is asserted by an utterance of that sentence, relative to those particular semantically relevant contextual features. Recanati argues that this characterization of semantic content will not help the propositional syncretist identify a “minimal” proposition, consistently with maintaining some of the positions characteristic of the Syncretic View.

Here is an example not discussed by Recanati, but which serves to make some of the same points he does. Many advocates of (either version of) the Syncretic View hold that quantifier domain restriction does not affect the proposition semantically expressed by a sentence relative to a context (though Soames is silent on this matter). If so, then the proposition semantically expressed by “Every bottle is in the fridge”, relative to any context, is the false proposition that every bottle in the entire universe is in the fridge. But this false proposition is *never* asserted or conveyed by an utterance of “Every bottle is in the fridge”. Thus, this characterization of semantic content is inconsistent with the view, advocated by so many adherents of the Syncretic View,

<sup>1</sup> Similarly, King and Stanley’s argument against sentence-level characters would be undermined if, as Schlenker (2003) has argued, there are operators that take sentence-level characters as objects.

<sup>2</sup> Recanati notes correctly that I have also provided a similar characterization of the semantic content of a sentence in context, in Stanley (2002). It never occurred to me that this kind of characterization of what is said could be used to defend the Syncretic View—it occurs in Stanley (2002) as a premise in a defense of a position antithetical to a number of the conclusions drawn in Soames (2002).

that domain restriction is semantically inert. The problem with using this version of the common denominator characterization of the minimal proposition is that it threatens to deliver the desired minimalist result only when “the asserted content is richer than the alleged semantic content” (Recanati (2004, 60)). But there are many cases (such as many cases of domain restriction) in which the semantic content, according to the advocate of the Syncretic View, is never itself asserted.

A natural response to these sorts of worries with the common denominator approach to the minimal proposition is to give up the view that the speaker must assert or intend to convey the minimal proposition, and seek some other relationship between the semantic content of a sentence relative to a context, and the information asserted by an utterance of that sentence. Indeed, Scott Soames, the target of much of chapter 4 of *Literal Meaning*, has recently done just this, abandoning his earlier conception of the relation between the two levels for reasons somewhat similar to the ones given above (Soames (2005)). In its place, Soames adopts a conception of the relation between the semantic content of a sentence in context and its asserted content that permits the semantic content of a sentence in context to be less than fully propositional (in which case, it is not a suitable candidate for a content to be asserted at all).

On Soames's new view (*ibid.*), the semantic content of a sentence in context may be a “propositional matrix”. The relation between the asserted content and the semantic content of a sentence in context is that the former must be “an acceptable completion” of the latter (*ibid.*, 365). This is an instance of non-propositional syncretism. Non-propositional syncretism, which allows semantic contents of sentences in contexts to be at least sometimes non-propositional entities, is close to Recanati's own position (2004, 56) that “semantic interpretation, characterized by its deductive character, does not deliver complete propositions; it delivers only semantic schemata—propositional functions, to use Russell's phrase”. Non-propositional syncretism and contextualism agree both in the thesis that semantic contents of sentences relative to contexts are often too underspecified to be propositional, and in the thesis that whatever is consciously available to the speaker as what she primarily intends to assert is not usually the semantic interpretation of the sentence uttered in that context. The two positions therefore have a great deal of similarity.

It is true that Recanati describes his position as one in which “the notion of ‘what a sentence says’ is incoherent”, suggesting that he thinks there is something wrong in principle with taking his non-propositional semantic schemata to be “what a sentence says”, whereas a non-propositional syncretist such as Kent Bach seems happy with taking propositional radicals to be what is said by a sentence. But this distinction is merely terminological. Another difference is that I suspect that Recanati believes that very few if any sentences, relative to contexts, have propositions as semantic contents (for example, Recanati (2004, 90)), whereas advocates of non-propositional syncretism

such as Kent Bach and Scott Soames think that quite a number of sentences have propositional semantic contents in context. However these positions are degrees upon the same continuum, and not (when one considers the degree of apparent contextual underspecification in language) very distant ones at that. The similarities between non-propositional syncretism and contextualism far outweigh their differences.

I have suggested that there is not a great deal of space between non-propositional syncretism and Recanati's favored version of contextualism. What, then, of propositional syncretism? Recanati's discussion of Soames, as well as Soames's subsequent advocacy of non-propositional syncretism, suggests that propositional syncretism is a difficult position to maintain. It is a matter of extreme difficulty to isolate the minimal proposition that is supposed to be the semantic content of a sentence in context. Cappelen and Lepore (2005) have provided an influential recent defense of propositional syncretism. Unfortunately, they do not say enough about the crucial question of how to isolate the minimal proposition. One suggestion they develop involves "[identifying] tests that help the theorist focus on the speech act content that a wide range of utterances of S have in common" (57). However, as Recanati and Soames (2005) emphasize, it is unlikely that the minimal proposition expressed will be something the speaker ever intends to assert, and so it is unlikely that the minimal proposition will be a part of *speech act* content. In another passage, Cappelen and Lepore suggest (143):

The semantic content of a sentence S is the content that all utterances of S share. It is the content that all utterances of S express no matter how different their contexts of utterance are. It is also the content that can be grasped and reported by someone who is ignorant about the relevant characteristics of the context in which an utterance of S took place.

This passage is rather unclear. Is it the content that all *assertions* of S express, no matter how different their contexts of utterance? If so, "every bottle is in the fridge" has no semantic content relative to any context, since there is one no proposition that is asserted by every utterance of the sentence (and certainly not, as we have seen, the proposition that every bottle in the universe is in the fridge, since this is *never* asserted). If the common content of all utterances of a certain sentence is not content that is the content of any genuine speech act, what is the motivation for thinking that common contents are always genuine propositions, rather than just Recanati's "semantic schemata"? After all, the main reason to think that the common contents are propositional is that they can be "claimed, asserted, questioned, investigated" (Cappelen and Lepore, 152). Once one sees that the common minimal contents are not the things claimed, asserted, questioned, or investigated, there is little motivation for believing them always to be propositions.

I share Recanati's skepticism that one can always isolate a common propositional content that can serve as the minimal content, consistently with other syncretic

commitments. Furthermore, even if the propositional syncretist were to assign common minimal propositional contents to each sentence in context via some artificial method, they would play no more a role in an account of communication than non-propositional entities such as semantic schemata. For such contents, though propositional, would no more plausibly be objects of intentional actions such as assertions, questions, or commands than schemata. So there is no important difference for the theory of meaning between propositional syncretism and non-propositional syncretism. I have already mentioned my suspicion that there is no great difference between non-propositional syncretism and Recanati's contextualism. My suspicion is therefore that the most important disputes in the theory of meaning are not between contextualism and syncretism, or even between advocates of these doctrines and semantic skeptics, such as Charles Travis. Rather, the genuinely important disputes in the theory of meaning are between those who maintain that the contents primarily asserted by speakers are *not* generally the semantic contents of the sentences they use (even relative to those contexts), and those who maintain that the contents primarily asserted by speakers *are* generally (not always, but typically) the semantic contents of the sentences used (relative to those contexts).

One worry with the position that speakers do not usually mean what their sentences express is that the intuitive distinction between literal meaning and non-literal meaning threatens to break down; the position seems to imply that we usually speak non-literally. In the fifth chapter, Recanati addresses this concern by arguing that there are two distinct literal/non-literal distinctions that are conflated by this worry. One can allow that speakers usually mean something different than the semantic content of the sentences they use, without endorsing the thesis that speakers usually mean something different than speakers in those situations would normally mean by the use of those sentences. By rejecting the "tacit assumption" that the semantic content of a sentence in context is what that sentence is usually used to say (81), the contextualist (or advocate of the syncretic view) frees herself from the problematic consequence that we usually speak non-literally.

The sixth chapter of *Literal Meaning* is devoted to setting up Recanati's arguments against the thesis he calls *indexicalism*. The semantic content of a sentence relative to a context is the result of combining the referential contents of the parts of that sentence, relative to that context, in accord with the composition rules determined by the syntactic structure of the sentence. *Literalism* is the implausible view that the intuitive truth-conditions of sentences relative to contexts are both the semantic contents of those sentences in the above sense, and determined entirely by rules of the language, independently of speaker's intentions. So for the literalist, the model for context-sensitivity is indexicals such as "I" and "today", whose denotation in context is fixed independently of speaker intentions. The indexicalist also holds that the contents primarily asserted by speakers are generally the semantic contents of



the sentences they use. Like the literalist, the indexicalist hypothesis is that the only way context affects the intuitive truth-conditions of an utterance is by helping to determine the interpretation of an element in the sentence used. In the vocabulary of Chapter 4, the indexicalist makes the empirical claim that there are no “strong pragmatic effects” on intuitive truth-conditions; in Recanati’s terminology, there are no “top-down influences”. Also like the literalist, the indexicalist rejects the existence of a substantial gap between the content a speaker primarily asserts when she utters a sentence, and the semantic content of that sentence, relative to that context of use. Unlike the literalist, however, the indexicalist recognizes that many context-sensitive expressions have their semantic contents fixed in part by reference to speaker intentions. Indexicalism is the plausible surviving descendant of literalism.

Why believe the indexicalist hypothesis? Recanati finds excluding “top-down” or “strong” pragmatic effects on intuitive truth-conditions “as dogmatic and stipulative as the literalist restriction of context-sensitivity to a short list of familiar indexical expressions” (160). But the indexicalist claim is an empirical hypothesis, not a stipulation about content. Furthermore, it is an empirical hypothesis that has a reasonable basis. There are special concerns with appeal to strong pragmatic effects. Linguistic communication is rule-governed and convention-bound in a way that would be mysterious, if there were strong pragmatic effects on intuitive truth-conditions.

There is both a very specific way to make this point, and a more general way. The specific way to make this point is via the problem of *over-generation* (Chapter 5). An utterance of (1) can be used to express what an utterance of (2) would have expressed. But an utterance of (1) cannot be used to express what an utterance of (3) expresses:

- (1) Every Frenchman is seated.
- (2) Every Frenchman in the classroom is seated.
- (3) Every Frenchman or Dutchman is seated.

The fact that (1) cannot be used to express (3) suggests that there is a conventional mechanism underlying the phenomenon of domain restriction. It is because of the conventional mechanism underlying domain restriction that certain sentences cannot be used to express certain propositions. In contrast, if interpreters had recourse to free pragmatic enrichment as a way of interpreting others, and speakers were aware that interpreters had recourse to this, then one should be able to use (1) to express (3). The fact that one cannot suggests that there are conventional linguistic mechanisms that govern permissible interpretations of the domain of quantified noun phrases.

If the literalist conception of language were correct, speaking would be an extremely effective means of communicating a particular message, but also rather unwieldy and impractical. It would take too long to find all the right words. But if the contextualist conception were correct, one might worry that speaking would be thoroughly

unnecessary for efficient communication. The indexicalist's position is a plausible starting hypothesis for how language is able to be sufficiently elastic as to be usable, and sufficiently rule-governed as to be useful. It is a reasonable empirical hypothesis, in advance of detailed inquiry.

Recanati (2004, 89–90) believes that the indexicalist starts off in a considerably weaker dialectical position than the contextualist:

Without going into the details, it is fair to say that the indexicalist starts with a significant disadvantage; for he makes a universal claim while his opponent only makes an existential claim. For his opponent to win, it is sufficient to produce *one* convincing example of a strong pragmatic effect. But the indexicalist is condemned to deal with *all* putative cases, and to show that they are not what they seem to be.

Recanati misstates the dialectical situation. The contextualist's method for arguing against indexicalism is to produce a reading R of an utterance of a sentence S, and argue that R is not the result of the semantic interpretation of S, relative to the context of utterance. In each such case, there are three different responses available to the indexicalist:

- (a) The first option is to establish that the alleged reading is not part of the intuitive truth-conditions of an utterance of that sentence, but is instead due to the pragmatics.
- (b) The second option is to argue that the claim that reading R is not due to the semantics is due to an overly simplistic conception of the semantic content of some elements of C. When the correct semantics for the relevant expression is given, the reading does emerge from the semantics.
- (c) The third option is to argue that the claim that reading R is not due to the semantics is due to an overly simplistic conception of the syntactic structure of C.

The advocate of strong pragmatic effects on intuitive truth-conditions must produce a case, and show that none of these options is available for that case. For each putative case in which it can be persuasively argued that (a) is not an option, the contextualist (or advocate of the syncretic view) must establish that there is *no way* of accounting for the problematic reading within the semantics. From this perspective, it is the contextualist who makes a universal claim. As Stephen Levinson (2000, 214) writes: "There will always be doubts about whether a better semantic analysis of the relevant construction might not accommodate the apparent pragmatic intrusions in some other way."

Establishing the existence of strong pragmatic effects on intuitive truth-conditions by appeal to particular examples is therefore no easy task. But Recanati does not rest his entire case on arguing that particular examples cannot be handled within

the semantics. In the beginning of chapter 7, he also provides a more theoretical argument for the existence of strong pragmatic effects, by appeal to his *Optionality Criterion* (2004, 101):

*Optionality Criterion*

Whenever a contextual ingredient of content is provided through a pragmatic process of the optional variety, we can imagine another possible context of utterance in which no such ingredient is provided yet the utterance expresses a complete proposition.

Recanati claims that the Optionality Criterion “gives us a criterion for telling apart cases in which a contextual ingredient results from saturation and cases in which it does not” (100). As an example, Recanati argues that the provision of a location to an utterance of “It’s raining” is a strong pragmatic effect. Recanati envisages a case (9–10) in which “rain has become extremely rare and important, and rain detectors have been disposed all over the territory (whatever the territory—possibly the whole Earth)”. When the rain detector goes off, and someone shouts “It’s raining!” what she expresses is that it is raining *somewhere or other*. According to Recanati, this is a case in which no location has been contextually provided. As Recanati comments (101):

Using the Optionality Criterion, however, I have established that the location of rain is not provided through saturation; for there are contexts in which the sentence “It is raining” expresses a complete proposition, even though no location is contextually provided as that which the utterance concerns.

Unfortunately, Recanati’s reasoning here is fallacious. The Optionality Criterion is a universally quantified conditional. It tells us that when we have a contextual ingredient of content that is provided through a pragmatic process of the optional variety, then Q is true, Q being that the utterance can be used to express a complete proposition even when no such ingredient is provided. But one cannot conclude from the fact that Q is true (that “it is raining” can be used to express a complete proposition, even though no location is contextually provided) that the antecedent is true, that the location is not provided through saturation. To do so would be to commit the fallacy of affirming the consequent. Yet that is just what Recanati infers from the Optionality Criterion, together with the premise that “It is raining” can be used to express a complete proposition, when no location is provided.

What Recanati requires for his argument, instead of the Optionality Criterion, is the following principle:

The Optionality Criterion\*

Whenever a contextual ingredient of content is provided, and we can imagine another possible context of utterance in which no such ingredient is provided yet the utterance expresses a complete proposition, then the contextual ingredient of content is provided by a pragmatic process of the optional variety.

The Optionality Criterion does not seem particularly objectionable; *if* an element has been provided by a pragmatic process of the optional variety, then it should not be necessary to provide it for the utterance to express a complete proposition. This principle is, however, utterly useless in helping to distinguish between cases in which a contextual ingredient is provided through saturation (a non-optional process) and when it is provided through an optional pragmatic process. What Recanati requires is the considerably stronger Optionality Criterion\*. But unlike the Optionality Criterion, there is no compelling reason to believe the Optionality Criterion\*.

There is no reason in principle to accept the Optionality Criterion\*. Suppose that there were unpronounced elements in the syntactic structure of a sentence that, like the overt English pronouns “he”, “she”, and “it”, could be used to refer to different contents on different occasions of use. Suppose that these elements, again like the overt English pronouns “he”, “she”, and “it”, could also be *bound* by unpronounced existential quantifiers, as the overt English pronouns “he”, “she”, and “it” can be bound by the overt quantifiers “someone” or “something”. In other words, suppose that in addition to overt pronouns and quantifiers, natural languages had a system of unpronounced pronouns and quantifiers. Presumably, it would be a language-specific matter which pronouns were pronounced; for example, in certain Romance languages (so called “pro-drop” languages), the pronominal subject of verbs do not need to be pronounced, though they are still syntactically active. This does not seem to be a particularly surprising or radical hypothesis. Indeed, it is the most conservative hypothesis governing the postulation of covert structure; that covert structure behaves like overt structure. Nevertheless, it is inconsistent with the Optionality Criterion\*. For in such a language, we would see contextual ingredients of content provided through a pragmatic process of the non-optional variety (saturation), yet utterances of the relevant overt sentences could be used to express complete propositions, even when no such ingredient is provided. In the latter kind of situation, the unpronounced pronoun would be existentially closed by an unpronounced existential operator.<sup>3</sup>

There is no *a priori* reason to accept the Optionality Criterion\*; there are easily conceivable accounts that are inconsistent with it. More worrisomely, there are widely accepted accounts of various phenomena that are inconsistent with it. For example, a consensus has developed in much of recent linguistics that tenses are not operators,

<sup>3</sup> In forthcoming work, Recanati raises the concern, presumably against a proposal of this sort, that overt pronouns are never bound by covert quantifiers. Well, this depends upon which analyses one regards as correct. On the analysis of Quantificational Variability Effects in Heim (1982), a sentence such as “A man walked in. He was wearing a hat” involves a covert existential quantifier binding both the open variable position in the indefinite “a man” and “he”. Similarly, some accounts of donkey anaphora treat the “it” in “Every man who owns a donkey beats it” as bound by a covert existential quantifier associated with the verb phrase.

but are rather predicates of times. In a sentence such as “John closed the door”, the past tense morpheme is best understood as a predicate of a contextually provided time, rather than as an operator. One piece of evidence for this conclusion is that tenses do not automatically iterate. For example, “Yesterday, John closed the door” does not mean that in the past of yesterday, John closed the door. “Yesterday” and the past tense morpheme behave like predicates of a uniform, contextually provided time (see King (2003) for a nice summary of the arguments for this conclusion).

The worry for the Optionality Criterion\* is that the position for the contextually provided time can also be bound. For example, the sentence “I shot a lion” can either be understood to express a proposition about a particular time (say, this morning), or be understood to express the existentially quantified proposition that at some time  $t$  in the past, I shot a lion at  $t$ . So, if the current consensus about tense is correct, then there is a pronominal element in the syntax of sentences that is assigned a time, relative to a context of use. But this pronominal element can also be existentially closed. So, when “I shot a lion” is used to express a proposition about a particular time, it is an instance of saturation. But “I shot a lion” can be used to express a complete proposition, even when no time is contextually provided. In such a case, the temporal variable is existentially closed. But, *contra* the Optionality Criterion\*, when a specific time is provided, it is still an instance of saturation.<sup>4</sup>

Furthermore, even if the Optionality Criterion\* were true, it would not help us in practice to distinguish contextual ingredients of content that were provided through saturation, and contextual ingredients of content that were provided through optional processes. As Recanati himself points out in the case of domain restriction (101–2), it is not clear whether we should take an unrestricted reading of an utterance of a quantified sentence to be a case in which the quantified expression *lacks* a contextually provided domain altogether, or rather take it to *have* a contextually provided domain, albeit one that is maximally large. For this reason, Recanati concludes that the Optionality Criterion (actually the Optionality Criterion\*) is of no use in deciding whether domains for quantifiers are provided via saturation or through a pragmatic process of the optional variety. It is curious that Recanati does not see that the very same point he makes about quantifier domain restriction also undermines his use of the Optionality Criterion\* to argue that the location for utterances of “It’s raining” is provided through a pragmatic process of the optional variety. For in Recanati’s envisaged example, it does seem like a location is provided. The location that is provided is *the whole earth* (not even the whole universe). So in order to make sense of his example, there does need to be a contextually provided location. Furthermore, even if rain detectors were set up everywhere in the universe, it is not clear whether

<sup>4</sup> In forthcoming work, Recanati has recognized the threat the standard account of tense described here poses for some of his views.

the relevant utterance of "It's raining" would have the whole universe supplied as the value of a contextual variable, or express a complete proposition in the absence of a contextually supplied location. In short, Recanati's correct argument that the Optionality Criterion\* is useless in deciding in practice whether the provision of domains for quantified expressions is the result of saturation or not applies *mutatis mutandis* to any argument using the Optionality Criterion\* that the location for an utterance of "It's raining" is provided through a pragmatic process of the optional variety.

Indeed, I suspect that the conclusions of the last paragraph can be generalized. That is, I am concerned that for every case in which there is a possible dispute between the indexicalist and the contextualist about whether a contextual ingredient is supplied through supplementation or a pragmatic process of the optional variety, similar reasoning will render the Optionality Criterion\* dialectically useless in deciding the matter. If so, then not only does the Optionality Criterion\* lack persuasive theoretical support, but even if we were to accept it, it would be useless in helping to adjudicate between the contextualist and the indexicalist.

In the rest of chapter 7, Recanati turns to a criterion advocated by indexicalists for detecting the existence of covert structure, which he calls the *binding criterion*:

#### The Binding Criterion

A contextual ingredient in the interpretation of a sentence *S* results from saturation if it can be "bound", that is, if it can be made to vary with the values introduced by some operator prefixed to *S*.

Using the binding criterion, indexicalists such as myself have argued that a variety of cases that contextualists have argued involve free pragmatic enrichment in fact involve saturation. Recanati sets out to undermine the binding criterion. First, he argues that the binding criterion would lead to an implausible multiplication of covert elements. Then, he provides a distinct semantic account of the relevant readings of the sentences in question.

Advocates of indexicalism have used the binding criterion in the following manner. Consider an utterance of the sentence (4), in New York City:

(4) It's raining.

Intuitively, this utterance expresses the proposition that it is raining in New York City. One might think that this is an instance in which the location of raining is provided by means of an "optional" pragmatic process, rather than via saturation. But consider:

(5) Whenever John lights a cigarette, it rains. (Chapter 1)

Intuitively, (5) expresses the proposition that whenever John lights a cigarette, it rains at the location at which John lights a cigarette. By the binding criterion, then,

the location is the result of the saturation of a covert pronominal element in the structure of (4).

Recanati's argument that the binding criterion leads to an implausible undermining of covert elements centrally involves the example:

- (6) Whenever John's father cooks mushrooms, John eats.

According to Recanati, "On a natural interpretation [of an utterance of this sentence], we understand that John eats *the mushrooms his father has cooked*." But in this example, "eats" is used intransitively. All participants in the debate agree that there is no argument place that can be bound by "whenever John's father cooks mushrooms" in the object position of "eats". So, the binding criterion predicts implausible covert structure.

However, as Luisa Marti (2006) has recently shown, Recanati's assumption that the intuitive interpretation of (6) involves binding is incorrect. (5) and (6) are not, despite appearances, analogous. Consider the discourse:

- (7) A: Whenever John's father cooks mushrooms, John eats.  
B: No he doesn't; curiously, he eats something else.

There is something decidedly odd about B's utterance; it certainly seems false. However, the discourse in (8) seems fine:

- (8) A: Whenever John lights a cigarette, it rains.  
B: No it doesn't; curiously, it rains somewhere else.

In other words, closer reflection on the content of (6) reveals that its natural interpretation is that whenever John's father cooks mushrooms, John eats *something*. So Recanati has not demonstrated that the binding criterion results in implausible multiplication of covert structure.

Recanati also provides a distinct account of the intuitive interpretation of sentences such as (5), one that does not involve postulating a variable for locations in sentences such as (4). Space considerations prevent me from providing a critique of Recanati's alternative account (which is in any case given in much greater detail in Recanati (2002)). Other contextualists and advocates of the Syncretic View have appealed to the pragmatic provision of variables to explain the relevant readings of sentences such as (5). But they have not bothered to explain how the process of pragmatically providing variables to sentences is to be constrained, leading to over-generation objections far more dire than any facing the indexicalist (Chapter 5). Recanati is unique among advocates of contextualism or the Syncretic View in attempting to provide an alternative positive account of the data that goes beyond such pragmatic magic tricks, and so this aspect of his work is to be particularly praised. If we are to move forward on our understanding of context-sensitivity in natural language, those who criticize the indexicalist position must follow Recanati's lead in assuming the obligation of

providing and defending alternative explanations of readings that are available, and those that are not.

Chapter 8 is devoted to the topic of circumstances of evaluation, in Kaplan's sense of circumstances with respect to which contents of sentences are evaluated. Recanati takes a special interest in this topic, because he sees a way of justifying his contextualist thesis that the semantic content of a sentence is something less than fully propositional. As Recanati (122) writes:

Once it is admitted that we need a circumstance over and above the content to be evaluated, we can part with Frege and, following Prior, tolerate contents that are not "semantically complete" in Frege's sense, that is, endowed with absolute truth-conditions. We can, because the circumstance is there which enables the content to be suitably completed. Thus the content of tensed sentences is semantically incomplete, yet the circumstance (the time) relative to which such a sentence is evaluated is sufficient to complete it.

The problem with Recanati's appeal to circumstances of evaluation to justify incomplete semantic contents is that it is in tension with much of current linguistic research. Most philosophers of language, and even many linguists, still accept that modals are operators of some kind (and so worlds are features of circumstances of evaluation). But, as I have indicated above, most linguists hold, *contra* Recanati, that tenses are not operators, and times are part of semantic content, rather than being features of circumstances of evaluation. Indeed, as King (2003) argues, the direction of research suggests that the only features of circumstances of evaluation are possible worlds.<sup>5</sup> Recanati must show this entire line of research to be incorrect. In particular, he must demonstrate the viability of (say) an operator account of phenomena such as sequence of tense, which have led researchers to treat tenses as predicates of times or events. This is a substantial obstacle to Recanati's program.<sup>6</sup> Indeed, one way of seeing the debate between indexicalism and contextualism is that the indexicalist position is the natural descendant of the trend in linguistic theory (starting with Partee (1973)) away from operator approaches of tense, and relativity of content generally, and towards explicit syntactic representation of elements that were once thought of as

<sup>5</sup> See Stanley (2005, ch.7), where this point is used against relativism about content generally.

<sup>6</sup> I am not claiming that it is impossible to provide an operator account of sequence of tense, in the sense of an account which, by clever manipulation of various covert indexed operators, generates the right truth-conditions. Rather, my claim is that such an account would be unwieldy and *ad hoc* (as King (2003, 219–21) points out). In particular, what results from such a framework is a proliferation of covert operators in the syntax that is considerably more cumbersome than the covert syntactic complexity which is anathema to Recanati and fellow contextualists and syncretists (see, e.g., Creswell (1996) for an example of the covert structure required to support an operator account of tense). So it is hard to see how someone with a general prejudice against postulating covert structure could *prefer* an operator account of tense to an account involving temporal variables.



features of circumstances of evaluation. If that is correct, a thorough-going defense of the contextualist position must establish that this trend is misguided. This is a task Recanati has yet to carry out.

In the last five years, a large amount of work has been devoted to assessing the success of giving a systematic semantic account of our intuitions about utterance content in the face of what James Higginbotham has called “the haze of use”. Most philosophers have coalesced around the position that our intuitions about utterance content float free of what is linguistically determined (with the natural result that there has been a great deal of internal bickering about terminology between those who occupy essentially the same position). In this book, Recanati does an extremely impressive job of laying out the interlocking commitments of his favored version of this position. Furthermore, he faces up to these commitments more honestly than almost any other member of the debate. Instead of merely objecting to other accounts, he gives his own analyses of various constructions, which then can be judged on their own merits. I have given some reasons to be pessimistic about the position he occupies, as well as some reasons to doubt the arguments he provides for it. Nevertheless, his book is essential reading for those interested in the debate.

## References

- CAPPELEN, HERMAN and LEPORE, ERNIE (2005), *Insensitive Semantics* (Oxford: Blackwell).
- CRESSWELL, MAX (1996), *Semantic Indexicality* (Kluwer).
- HEIM, IRENE (1982), *The Semantics of Definite and Indefinite Noun Phrases* (Amherst: University of Massachusetts dissertation).
- KING, JEFFREY (2003), ‘Tense, Modality, and Semantic Value’, in J. Hawthorne and D. Zimmerman (eds.), *Language and Philosophical Linguistics, Philosophical Perspectives*, 17 (Oxford: Blackwell).
- LEVINSON, STEPHEN (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicatures* (Cambridge: MIT Press).
- MARTI, LUISA (2006), ‘Unarticulated Constituents Revisited’, *Linguistics and Philosophy*, 29(2): 135–66.
- PARTEE, BARBARA (1973), ‘Some Structural Analogies between Tenses and Pronouns in English’, *Journal of Philosophy*, 70: 601–9.
- RECANATI, FRANÇOIS (1993), *Direct Reference: From Language to Thought* (Oxford: Blackwell).
- RECANATI, FRANÇOIS (2002), ‘Unarticulated Constituents’, *Linguistics and Philosophy*, 25: 299–345.
- SCHLENKER, PHILIPPE (2003), ‘A Plea for Monsters’, *Linguistics and Philosophy*, 26: 29–120.
- SOAMES, SCOTT (2002), *Beyond Rigidity* (Oxford: Oxford University Press).

- SOAMES, SCOTT (2005), 'Naming and Asserting', in Z. Szabo (ed.), *Semantics Versus Pragmatics* (Oxford: Oxford University Press), 356–82.
- STANLEY, JASON (2002), 'Modality and What is Said', in J. Tomberlin (ed.), *Philosophical Perspectives, 16: Language and Mind* (Blackwell), 321–44.
- STANLEY, JASON (2005), *Knowledge and Practical Interests* (Oxford: Oxford University Press).

# Postscript

The chapters in this book were written over a seven year period, from 1998 to 2005. Naturally enough, some changes in view occurred during their writing. There were some unclarities in the original papers, and some arguments against my positions that I failed to address. My purpose in this postscript is to explain the current “state of play” with regard to the issues discussed in these chapters. Since I have a number of distinct points to discuss, I will divide them into sections.

## I The Location of Domain Indices

The first point involves not a change of view but a final resolution of my indecision between two distinct views. It involves the account of quantifier domain restriction defended in “On Quantifier Domain Restriction” and “Nominal Restriction”. In these two chapters, the account is presented as a proposal in lexical semantics. Nouns are said to co-habit terminal nodes with domain indices. In “Context and Logical Form”, by contrast, no mention is made of this proposal. This is because the idea that the theory should take the form of a proposal in lexical semantics was (I think) mainly Szabó’s in “On Quantifier Domain Restriction”. I was neutral about it, and while writing “Context and Logical Form”, did not consider it as an option. However, by the time I wrote “Nominal Restriction”, I had come around completely to the view that Szabó had been urging on me. In particular, I was impressed by the fact that adoption of the lexical semantic perspective (in effect, treating each nominal as a compound noun of an articulated part and some domain variables) helped to evade certain objections to the proposal.

However, in due course I was brought back to the view in “Context and Logical Form” that domain indices should occupy their own terminal nodes (and, indeed, this is the view I defend in “Semantics in Context”). Three considerations led to

this decision. First, the arguments in favor of domain variables are syntactic in nature. They support the existence of pronominal elements in the logical form of sentences containing quantified noun phrases. They are not considerations from lexical semantics. Evidence for syntactic reality is evidence from considerations provided in support of genuine syntactic structure. This point was clear to me in “Context and Logical Form”, where I provided several arguments in favor of the syntactic reality of domain variables, including “strict–sloppy” ambiguities. Secondly, the arguments against the syntactic reality of domain variables that had led me to accept the lexical semantic view of domain restriction, which involve the difficulty of anaphora on such variables, are independently unconvincing (as I argue in a footnote in “Semantics in Context”). In particular, the syntactic reality of many elements less controversial than the existence of domain variables would be threatened by them. Third, certain arguments by Richard Breheny (2004) and Timothy Williamson (2004) showed that adjectives sometimes take scope over the nominal, without taking scope over domain variables associated with the nominal. Suppose we are at a costume party where certain people are pretending to be philosophers. I say:

(1) Every fake philosopher has a moustache.

If the domain restrictor were part of the noun “philosopher”, then “fake” would take scope over “philosopher”. But an utterance of (1) does not express the proposition that everyone who is faking having the property of being a philosopher and being at the party has a moustache, for this would include people who are faking being at the party. So, “fake” does not take scope over the domain variables. This suggests that the domain variables associated with “philosopher” must be attached to the adjective–noun pair “fake philosopher” and so cannot co-habit a node with the noun. If domain variables occupy their own terminal node, and are syntactically like relative clauses associated with nominals, then they could adjoin to the adjective–noun combination. So, while I have no final stated view about the proper syntactic representation of domain variables, as my considerations are neutral between several different hypotheses, I no longer think that domain variables co-habit terminal nodes with nouns.<sup>1</sup>

<sup>1</sup> There is a way to respond to the Breheny–Williamson worry. “Fake” is an intensional adjective. Perhaps the domain is rigidified along the appropriate intensional dimension. This would achieve the purpose of wide-scope, without implementing it syntactically. As Williamson (p.c.) has pointed out to me, the issue also arises with a sentence such as “Every non-philosopher has a moustache”, where “non” does not have scope over the domain variables associated with “philosopher” (and of course “non” is perfectly extensional). But here one can just have the domain variables attach to the compound “non-philosopher”.

## II Variable-rich Representations versus Variable-free Representations

My project is chiefly directed not at linguists but at those philosophers and cognitive scientists who have argued that semantic content was vastly impoverished relative to the intuitive truth-conditions of an utterance. My purpose was to show these theorists that the arguments they employed in favor of their conclusions under-estimated the resources available to the semanticist. However, in building a detailed case, I inevitably incurred a number of commitments on the syntax and semantics of the constructions I discussed. For example, I advocate “variable rich” syntactic representations to account for behavior inside a phrase that seems pronominal in nature. Advocates of variable free semantics, such as Pauline Jacobson, reject this approach, choosing instead to account for bound readings semantically (via type-shifting). Thus, an advocate of variable free semantics rejects the “binding assumption” discussed in “Context and Logical Form”. Within variable-rich approaches, my detailed proposals also incur quite specific commitments. For instance, Szabó and I argue that domain restriction for quantified noun phrases is associated with the noun rather than the determiner, and I employ this theory to account for a variety of apparently unrelated phenomena in “Nominal Restriction”. Many linguists broadly sympathetic with my general approach differ about these specific commitments. It is important to distinguish the general question of how much of what is said is linguistically determined from more detailed issues about the manner in which the linguistic determination of intuitive content comes about.

For example, variable-free semantics is a distinct approach to the binding data I discuss in many of the chapters in this volume. But it would be a mistake to think that appeal to variable-free semantics provides any succor at all to the advocate of free pragmatic enrichment. As I briefly explain in “Semantics in Context”, a variable-free framework such as that advocated in Jacobson (1999) does not allow for free pragmatic enrichment. What would be captured by the postulation of an unpronounced free variable, in a variable-rich framework like my own, would be, on Jacobson’s theory, captured by means of a type-shifting rule in the semantics. On my approach, a sentence such as “John failed three students” expresses a full proposition, one that results from contextual saturation of free variables associated with the nominal “students”. On Jacobson’s theory, “John failed three students” ends up expressing a function from the type of quantifier domains to propositions. Jacobson’s approach therefore does not admit free enrichment; the provision of a property as a quantifier domain is linguistically controlled, and is so in the same manner as the provision of a property to a contextual index would be. That is, on my account, domain restriction is linguistically controlled, since the domain has to match the linguistic type and meaning of free variables that determine it in the syntax. On Jacobson’s account,

domain restriction is linguistically controlled, since the domain provided by context has to match the type called for by the semantic value of the sentence in that context.

So, disputes about the details of my proposals take us away from the main issue, which is whether there is a plausible way to defend the thesis that we can capture communicated content entirely by means of consciously accessible, linguistically controlled content (content that results from semantic value together with provision of values to free variables in the syntax, or semantic value together with the provision of arguments to functions from certain semantic types to propositions) together with general conversational norms. For this reason, I tried to stay as neutral as possible about specific implementations in “Context and Logical Form”.<sup>2</sup>

### III Adverbs of Quantification

One worry that I have often heard involves adverbs of quantification, such as “always”, “sometimes”, or “usually”. Such expressions also have restricted readings, as in “A cat usually lands on its feet”. But the theory of domain restriction introduced in

<sup>2</sup> Pauline Jacobson (2006) has raised the worry, about my discussion of Breheny’s examples in “Semantics in Context”, that there are related examples that raise the same point, i.e., show the need for multiple domain variables for a single quantified noun phrase, but for which my response is inadequate. Her example involves a multi-disciplinary party at which there are philosophers of language, linguists, and psychologists (I have altered the example somewhat). Eager to place their semantics graduate students into valuable philosophy of language jobs, the linguists decide to poison the philosophers. Since linguists are not known for their ability to mix drinks, they practice first mixing non-lethal drinks for the psychologists; each linguist is pre-assigned a unique philosopher and a unique psychologist for whom to pour drinks. The plan worked because:

- (1) No linguist told any philosopher that the drink was poisoned.

Jacobson’s concern is that there is no way, with the use of a single function and argument, to obtain the intuitive reading of (1), which is that no linguist told any philosopher that *the drink that linguist poured for that philosopher* was poisoned. According to Jacobson, obtaining the intuitive reading of (1) requires two argument variables, one bound by the quantifier “no linguist”, and the other bound by the quantifier “no philosopher”. But on my account, the quantified noun phrase “the drink” is associated only with a single argument variable. But my framework does account for the intuitive reading of (1). According to my theory, the structure of the quantified noun phrase “the drink” is “the drink  $f(i)$ ”, where “ $f$ ” is assigned a contextually salient function, and “ $i$ ” is an argument variable. The story Jacobson tells raises to salience a function that, for each linguist  $A$  as argument, yields the property:

- (2)  $\lambda x(x$  is poured by  $A$  for  $ty(\text{philosopher}(y) \ \& \ \text{Assigned-to}(y, A))$ )

In short, the function raised by salience by her story is a function mapping linguists to the property of being a drink that linguist pours for the philosopher assigned to her. This smoothly accounts for the intuitive reading of (1), without the need to postulate an additional argument variable.

“On Quantifier Domain Restriction” and developed in “Nominal Restriction” does not apply to domain restriction for adverbs of quantification, such as “always”, “sometimes”, and “usually”. For according to the Nominal Restriction Theory, domain restriction is associated with nouns rather than quantificational determiners such as “every” or “some”. But adverbs of quantification are not accompanied by nouns. So, the theory does not generalize to cover domain restriction for adverbs of quantification, and thus renders domain restriction some kind of non-natural kind.

The point that the nominal restriction theory does not extend to domain restriction for adverbs of quantification is usually made in support of theories that associate domain indices with the quantificational determiner, rather than the noun. But I have difficulty seeing how it could support such theories. Quantificational determiners are of a different syntactic and semantic category than adverbs of quantification. So associating domain indices with quantificational determiners no more helps to explain how domain restriction with adverbs of quantification functions than does the Nominal Restriction Theory. According to both the nominal restriction theory, as well as theories that associate domain variables with quantificational determiners, there is a mechanism that explains how domain restriction is linguistically controlled for quantified noun phrases, and that mechanism is not relevant for explaining how domain restriction is linguistically controlled for adverbs of quantification. If this makes domain restriction into a non-natural kind, so be it.

The existence of a distinct mechanism for restricting the domain of adverbs of quantification undermines some recent critiques of certain applications of the Nominal Restriction Theory. Both Berit Brogaard (forthcoming) and Delia Graff Fara (forthcoming) have argued that my discussion of quantificational variability effects with definite descriptions in the conclusion of “Nominal Restriction” is somewhat problematic. In that paper, I consider the sentence:

- (2) [The <customer,  $f(i)$ >] is always<sub>i</sub> right.

I argue that the intuitive reading of (2) is derived via binding of the domain indices associated with the noun “customer” by the adverb of quantification “always”. But Brogaard and Graff Fara point out that, if “always” is unrestricted, that (2) will be false or undefined, since the contextually salient function that is the value of “ $f$ ” will be undefined for many values of “ $i$ ”. But I assume that the distinct mechanism that accounts for domain restriction for adverbs of quantification such as “always” will limit the quantification over situations to situations for which the salient function is defined. This is not to say that I am completely happy with my brief discussion of quantificational variability effects in the conclusion of “Nominal Restriction”. If what I say there is right, there should be quantificational variability effects present with all quantified noun phrases, since the phenomenon is due to an adverb of quantification binding a domain variable associated with a noun. As Donka Farkas and Delia Graff Fara

(forthcoming) have recently emphasized, however, some determiners do not easily allow quantificational variability effects at all, as witnessed by the contrast between:

- (3) A policeman usually eats donuts.
- (4) Some policeman usually eats donuts.

An utterance of (3) naturally expresses the proposition that most policemen eat donuts; (4) does not permit this reading. But, given my approach to quantificational variability effects, this distinction is utterly mysterious. Clearly, the acceptability of quantificational variability effects is due to the determiner rather than the noun. This is not to say that many quantificational variability effects are due to the binding of domain indices; it is just to say that more needs to be said about why certain determiners, such as “some”, do not allow such readings.<sup>3</sup>

## IV Adjectives

Another point related to the chapter “Nominal Restriction” concerns my treatment of adjectives in section II of that chapter. As Michael Glanzberg (p.c.) has emphasized, it is not clear which of two positions I am advocating about the underlying syntactic structure of predicative uses of adjectives, such as the use of “tall” in “John is tall”. On the first interpretation, I am advocating the view that “tall” is associated with an unpronounced comparison class nominal, which is in turn associated with a domain restrictor of the form “f(i)”. On the second interpretation, I am advocating the view that “tall” is associated simply with a domain restrictor of the form “f(i)”, which determines the comparison class for that adjective. Considerations of parsimony pull in two directions here. On the one hand, if adjectives are associated with unpronounced nouns, then the account is just an application of the theory of Nominal Domain Restriction. On the other hand, if “tall” is just associated with a domain restrictor of the form “f(i)”, which determines a comparison class, then correct interpretation does not involve determining a property, a function, and an object intended by the speaker (or similar properties, functions, and objects), but merely a function and an object. Since the chapter is a defense of the Nominal Restriction Theory, it was however the first interpretation I had in mind. I suspect that the unpronounced nominal occasionally simply has a vacuous content (say, the meaning of the noun “thing”, which applies to everything). This would resolve the additional interpretive pressure on hearers. In other cases, perhaps the function variable yields a function from an

<sup>3</sup> If Farkas (2002) is correct, this has to do with the differing discourse properties associated with “some” and “a”.



object to the class or plurality of all things, in which case the value of unpronounced nominal will be doing all the work of providing the comparison class.

So here is the basic proposal, as applied to predicative uses of adjectives. The logical form of (5) is roughly as in (6):

- (5) John is tall.  
 (6) John  $\lambda i(\text{is tall } N f(i))$ .

The variable “N” ranges over possible noun denotations (i.e. properties), the variable “f” ranges over functions, and the variable “i” ranges over objects. Relative to a context in which fourth-graders are at issue, the noun variable “N” would be assigned the property of being a fourth-grader, and “f” would denote a function from objects to the class or plurality of all things. Relative to such a context, (6) would express the proposition that John is tall for a fourth-grader.

Thus, on this proposal, there is not much of a structural difference between predicative uses of adjectives and uses of adjectives that occur with explicit nominal complements. On my proposal, the syntactic structure of (7) is in (8):

- (7) John built a tall snowman.  
 (8) John  $\lambda i(\text{built a tall snowman } f(i))$ .

The difference between a predicative use of an adjective and a use of an adjective with an explicit nominal is only that the latter occur with explicit determiners. I am inclined towards the view that a predicative use of an adjective occurs with an empty indefinite determiner, with the same meaning as the English indefinite determiner “a”. That makes the difference between (5) and (7) even more minimal.

One point that did not occur to me when writing “Nominal Restriction” was that Hans Kamp’s example (7) is just another way to make the argument due to Delia Graff Fara that shows that domain restrictors must be associated with the nominal, rather than the determiner. Recall that in Kamp’s envisaged example, (7) is used to express the proposition that John built a tall snowman *for a snowman built by a 7-year-old*. Kamp’s point, as represented in my framework, is that the adjective “tall” must take scope over the noun “snowman” as well as its associated domain restrictor, which determines the property of being a snowman built by a 7-year-old (I assume that “f”, relative to the envisaged context, is assigned a function that takes John to the property of being built by someone of John’s age). So Kamp’s example, treated in my framework, also shows that adjectives must take scope over domain restrictors.

The resulting theory of adjectives accounts for a number of facts about adjectives. First, as Ewan Klein (1981) noted, comparison classes for adjectives undergo “strict–sloppy” ambiguities similar to the ones apparent with explicit pronouns, as in the ambiguity (9) has between (10) and (11):

- (9) John loves his mother and Bill does too.

- (10) John loves John's mother and Bill loves Bill's mother. (sloppy)  
 (11) John loves John's mother and Bill loves John's mother. (strict)

The standard account of this ambiguity is that it results from the distinction between bound and referential uses of the pronoun "his".<sup>4</sup> In particular, the "sloppy" reading (10) of (9) is one on which the syntactic structure of (9) is:

- (12) John  $\lambda x$ (loves  $x$ 's mother) and Bill does too.

The ellipsis "does too" is elliptical for the Verb Phrase " $\lambda x$ (loves  $x$ 's mother)", and on this reading, (9) yields the interpretation as in (10). The sort of evidence discussed by Klein is captured similarly, on my theory of adjectives. As Klein points out, (13) can have the reading given in (14):

- (13) That elephant is large, and that flea is too.  
 (14) That elephant is large for an elephant, and that flea is large for a flea.

On my account, reading (14) of (13) is due to the fact that the underlying syntactic structure of (13) is:

- (15) That elephant  $\lambda i$ (is large N  $f(i)$ ) and that flea is too.

Relative to the envisaged context, the noun "N" is assigned the property of being a thing (which is true of everything), and "F" is assigned a function from entities to the species to which they belong. At logical form, (15) is then:

- (16) That elephant  $\lambda i$ (is large N  $f(i)$ ) and that flea is  $\lambda i$ (is large N  $f(i)$ ).

Relative to the envisaged context, (16) expresses the same proposition as (14). So, my theory straightforwardly predicts the existence of strict-sloppy ambiguities for adjectives in verb phrase ellipsis.

Finally, as I emphasize in section IV of "Semantics in Context", my treatment also accounts for some particularly subtle bound readings of comparative adjective variables, as in:

- (17) Most sports teams have members that are old.

On my account, the relevant syntactic structure of (17) is given in (18):

- (18) Most sports teams  $\lambda i$ (have members that are old N  $f(i)$ ).

<sup>4</sup> Nothing is unrevisable (or indeed unrevised) in linguistic theory, but this account of the strict-sloppy ambiguity has survived many challenges over the years. For example, Paul Elbourne (2006) uses the phrase "the common view" for the conjunction of three claims, one of which is this account of the strict-sloppy ambiguity (Elbourne seeks to defend the common view). But it is worth mentioning that this account of the strict-sloppy ambiguity is perhaps the least problematic of the three theses that constitute Elbourne's "common view".

Relative to the envisaged context, “N” is again assigned a vacuous property denotation, and “f” is assigned a function that takes a sports team to the property of being someone who plays that sport, or perhaps being someone who plays that sport well. Suppose that the contextually supplied function (for example, the one intended by the speaker) is a function from sports teams to participants in that sport; i.e. it takes a sports team, and yields the property of being a person who plays that sport. Then the account predicts that (18) expresses the proposition that most sports teams have members that are old relative to the class of people who play that sport. Even if every sports team has members who play more than one sport, this treatment of (17) yields the correct intuitive truth-conditions. I argue in section IV of “Semantics in Context” that it is not clear what other accounts would have similar success.

John Hawthorne (p.c.) has raised the concern that there are readings of adjectives where the intuitive comparison class differs from the nominal restriction. For example, suppose I say, of a man on the North Pole:

(19) He is a remarkable violinist.

By “remarkable” one may mean “remarkable by the standards of violinists in Europe”. But since the man is not a violinist in Europe, it is difficult to see how (19) could be true. The response is that the domain for “remarkable violinist” is not the property of being a violinist in Europe. It is rather the property of being a violinist in Europe or being the demonstrated man. This property is not a natural one. But, as I emphasize in “Nominal Restriction”, the properties that are comparison classes do not need to be natural kinds.

One might wonder how my treatment of adjectives relates to the degree theoretic treatment of adjectives that has recently achieved the status of generally accepted theory, thanks in large part to the work of Kennedy (1999). The answer is that my treatment of adjectives is fully consistent with degree theoretic approaches. My discussions of adjectives do not amount to a theory of the semantics of adjectives; I say nothing about the meaning of “old” or “tall” in the above examples. Rather, my discussions are a contribution to the theory of domain restriction, and derivatively to an account of how context supplies comparison classes for adjectives. According to my proposal, predicative uses of adjectives in fact involve underlying structures similar to adjectives with noun complements. The degree theorist about adjectives has a theory about the semantics of sentences such as (6), and if the theory works for such sentences, it will also work for predicative uses of adjectives.

According to my theory of adjectives, all the context-sensitivity associated with comparison classes can be traced to a nominal restriction index. This is a bold claim, and one that may very well turn out to be excessively restrictive. Perhaps each adjective must be associated with its own comparison class. But I prefer to advance the more

restrictive claim, and see where it may fail, than to adopt a more permissive claim, and add on *ad hoc* restrictions.

## V “It’s raining”

There has been some confusion about my discussion of “It’s raining”, the classic example of unarticulated constituents, in section III of Chapter 1. Recall that the argument for unarticulated constituents is, very roughly, as follows. An utterance of “It’s raining” communicates a proposition about a specific location; say, New York City. Yet there is no variable for locations in the syntactic structure of the sentence “It’s raining”. So, the location is an unarticulated constituent of the proposition expressed by an utterance of “It’s raining”.<sup>5</sup> The purpose of my discussion of this sort of example in “Context and Logical Form” was to show that one could account for this sort of example, without postulating unarticulated constituents of utterances.

My discussion involved the following sort of example:

(20) Every time John lights a cigarette, it rains.

In order to capture the intuitively natural reading of (20), the quantifier “every time” must bind some sort of variable in the sentence “it rains”. I proposed two sorts of accounts of how this could work (p. 53). According to the first, one can “replace the assumption that ‘rain’ introduces a hidden temporal variable with the assumption that it introduces a hidden situation or event variable, which can either be bound . . . or free”. According to the second, in the logical form of “it rains”, there are two variables, one yielding a location and the other yielding a time. It is in fact the first of these proposals that I prefer. However, the literature that has responded to my discussion (for example, Recanati 2004, ch. 7) has only focused on the second, less plausible, account. So I will explain how the first suggestion, involving event variables, explains the data without the postulation of unarticulated constituents.

On the envisaged framework, instead of being associated with temporal variables, verbs are associated with event or situation variables. These event or situation variables

<sup>5</sup> Sennet (2006) discusses some difficulties surrounding the precise definition of an unarticulated constituent. As he points out, it is not sufficient to say that an unarticulated constituent is an element of the proposition expressed that is not the value of any expression in the sentence uttered. For example, an utterance of “The mayor of New York City said that it is raining” contains an expression that refers to New York City, but New York City is still an unarticulated constituent of the utterance, according to fans of unarticulated constituents. The reason it is still an unarticulated constituent is because the structured proposition expressed by this utterance contains more occurrences of New York City than the sentence contains expressions that denote it.

can either be bound or free. In a bare use of “It’s raining”, the speaker makes deictic reference to a particular event or situation, and says of it that it is a raining event. The event has a location (say, New York City). So, on this account, an utterance of “It’s raining” is about a particular event, which is the value of an event variable in the syntax of the sentence. The location is not, on this account, an unarticulated constituent of the utterance; rather the proposition expressed by the utterance contains an event, and the speaker and hearer know where the event is taking place because of their general knowledge about the world.

On this framework, event or situation variables take the place of temporal variables. The logical form of (20) would then be:

$$(21) \forall e(\text{John lights a cigarette at } e \rightarrow \text{rains}(e))$$

Some other examples involve slightly distinct treatment. Recanati (2004, 104) suggests the example:

$$(22) \text{Everywhere I go, it rains.}$$

Clearly, “everywhere” quantifies over locations rather than events. Possible representations of (22), on the event account, include:

$$(23) \forall x \forall e (e \text{ is a going to } x \text{ by me} \rightarrow (\text{rains}(e))).$$

$$(24) \forall x (\text{I go to } x \rightarrow \exists e(\text{rains}(e) \text{ and } \text{At}(x,e))).$$

On all of these accounts, one can account for the relevant bound readings without postulating unarticulated constituents for bare uses of “It’s raining”.

## VI Non-Sentential Assertion

In section II of “Context and Logical Form”, I argue against the existence of non-sentential assertion. Philosophers have found my contention that there is no non-sentential assertion particularly difficult to accept, and as a result have not carefully read the arguments I propounded. I want briefly to return to that discussion, to correct some misimpressions. It is worth emphasizing that there has also been a growing literature in syntax developing the thesis I defend in this section. The arguments given in this literature support bolder and more dramatic conclusions than the one I sought in this section, and are also (quite frankly) supported by better arguments, some of which I briefly discuss below.

In “Context and Logical Form”, I pursued what Robert Stainton has called a “divide and conquer” strategy (see Stainton (2006) for a vigorous defense of the existence of non-sentential assertion). First, I argued that a number of apparent

cases of non-sentential assertion were in fact assertions after all, because they were elliptical on contextually salient questions. But I conceded, in this section, that there were cases that could not be handled in this manner. So, secondly, I argued, in the case of these examples, that they were not in fact assertions. My argument that genuine (non-elliptical) non-sentential speech did not amount to assertions employed a premise about genuine assertions, which was that genuine assertions express unique propositions “up to vagueness”. My intent by the latter caveat was to include approaches to vagueness or indeterminacy or context-sensitivity that involve assertions expressing multiple propositions. For example, according to William Blackburn (1998), an utterance of say “Every bottle is in the fridge” expresses that set of propositions of the form “every bottle that is F is in the fridge” which are consistent with the speaker’s intentions. One might also have an approach to vagueness that treats a vague expression as indeterminately expressing many distinct yet related properties. On this approach, an utterance of “That is a heap” expresses a set of propositions, which differ from one another in containing distinct yet related possible denotations of “heap”. My premise is consistent with such approaches to domain restriction and vagueness. But it is not consistent with views of assertion according to which an utterance asserts multiple propositions, some of which are conversational implicatures of others. This is a view that has become increasingly popular in philosophy of language. My argument that utterances of genuinely sub-sentential speech are not assertions has no force for proponents of this liberal conception of assertion.

However, I have been convinced by the subsequent literature on the topic that it was a mistake to pursue the “divide and conquer” strategy. I should simply have argued that there are no genuine cases of complete non-sentential utterances. My strategy was to show that a number of cases of non-sentential assertion actually seemed only to be felicitous in contexts in which there was a contextually salient question, one which served as an antecedent to verb phrase ellipsis. But a number of other arguments have emerged in the literature to support the conclusion that there is no non-sentential assertion. First, as Peter Ludlow (2005) has pointed out, “often the sub-clausal fragment that is actually pronounced could not be generated unless it was the product of clause-level operations”. In short, many cases of apparent non-sentential speech could not be generated, according to standard assumptions, unless there was clausal (i.e., sentential) structure. For example, one old account of passives is that they are derived via movement from non-passive structures; as Ludlow shows, clausal structure would be required to support such an analysis. So, classic cases of apparent non-sentential speech such as “Hood sunk” would be sentential after all. A related argument, pursued both in Ludlow (2005) and in Merchant (2004), is that apparent non-sentential speech involves fragments that bear the kind of case-marking they would have if they were part of a larger sentence. There are different accounts

of case-marking in different syntactic frameworks. But common to most accounts is that case-marking involves something like clausal structure. For instance, on basic Minimalist approaches, dative case-marking involves movement to check case features through clausal structure, and on the Extended Standard Theory it involves assignment of case by a verbal element. In sum, many considerations have emerged in the literature to support the conclusion that genuine non-sentential speech does not naturally occur. I concur with the critique given of my discussion in Merchant (2004), namely that I was too quick to concede that many cases of apparent non-sentential assertion are genuinely non-sentential.

## References

- BLACKBURN, W. K. (1998), 'Wettstein on Definite Descriptions', *Philosophical Studies*, 53: 263–78.
- BREHENY, RICHARD (2004), 'A Lexical Account of Implicit (Bound) Contextual Dependence', in *Semantics and Linguistic Theory* (SALT), 13.
- BROGAARD, BERIT (forthcoming), 'Descriptions: Predicates or Quantifiers?', in *Australasian Journal of Philosophy*.
- ELBOURNE, PAUL (2006), 'Split Antecedents', Semantics Archive.net.
- FARKAS, DONKA (2002), 'Varieties of Indefinites', Proceedings of SALT, 12.
- GRAFF FARA, DELIA (forthcoming), 'Descriptions with Adverbs of Quantification', *Philosophical Issues*.
- JACOBSON, PAULINE (1999), 'Variable-free Semantics', *Linguistics and Philosophy*, 22: 117–84.
- JACOBSON, PAULINE (2006), 'Direct Compositionality and Variable-free Semantics: Taking the Surprise out of "Complex Variables"', handout from the 30th Penn Linguistics Colloquium.
- KENNEDY, CHRIS (1999), *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison* (New York: Garland Press).
- KLEIN, EWAN (1981), 'A Semantics for Positive and Comparative Adjectives', *Linguistics and Philosophy*, 4: 1–45.
- LUDLOW, PETER (2005), 'A Note on Alleged Cases of Non-sentential Speech,' in R. Elugardo and R. Stainton (eds.), *Ellipsis and Non-Sentential Speech*, Studies in Linguistics and Philosophy, 81 (Dordrecht: Springer).
- MERCHANT, JASON (2004), 'Fragments and Ellipsis', *Linguistics and Philosophy*, 27/6: 661–738.
- SENNET, ADAM (2006), 'Argument, Adjuncts, Slopes and Binding: or What is Not Said Aloud in Polite Conversation', diss. Rutgers University.
- STAINTON, ROBERT (2006), *Words and Thoughts* (Oxford: Oxford University Press).
- WILLIAMSON, TIMOTHY (2004), 'Everything', in *Philosophical Perspectives*, 26: *Philosophical Linguistics and Philosophy of Language*.

# INDEX

- adjective 11, 23–4, 26–7, 19, 223 n.15, 249, 253–6  
     comparative 54–5, 83–4, 112, 122–4, 127 n.11, 192–4, 210–11, 216–7, 220 n.12  
     gradable 20, 203  
     superlative 119 n.6  
 adverb  
     of quantification 49 n.21, 130–1, 251–2  
     temporal 164–5  
 Ajdukiewicz, Kazimierz 49 n.23  
 ambiguity 7, 11, 19, 37–9, 74, 76 n.9, 77 n.10, 78, 82 n.16, 129, 140 n.7, n.9, 158 n.32, 165, 187–8  
     collective/distributive 175, 203  
     deictic/bound 54 n.31, 255  
     intersective/non-intersective 124 n.10  
     strict/sloppy 60–2, 249, 254–5  
 anaphora  
     cross-sentential 106–7  
     and deferred reference 206  
     and domain variables 113–14, 220 n.12, 249  
     donkey 241 n.3  
     E-type 42 n.15  
 Austin, J. L. 31, 37  
  
 Bach, Kent 10, 24, 31, 32 n.4, 35–6, 47 n.20, 54, 56 n.35, 57, 74 n.6, 78 n.13, 85, 86 n.20, 90–93, 95 n.31, 125, 135, 142 n.16, 158, 159 n.33, 183, 186, 191–2, 195, 199, 200, 203, 215 n.9, 233–6  
 Bar-Hillel, Y. 32  
 Barton, E. 39, 41  
 Barwise, J. 96 n.33, 102, 115  
 Bezuidenhout, Anne 13  
 binding argument 23, 27, 48–58, 92–5, 100–101, 114, 119, 127, 165, 184–6, 191–7, 212–16, 218, 220–7, 257  
 binding assumption 49–51, 250  
 binding criterion 225, 243–4  
 binding, pragmatic 192–5, 199–200, 220  
 Bittner, Maria 54 n.31  
 Blackburn, W.K. 45 n.18, 87, 259  
 Braun, David 72 n.2, 73 n.3, 99 n.39, 138 n.4, 145–7, 155 n.30  
 Breheny, Richard 222–5, 249, 251  
 Brogaard, Berit 252  
 Borg, Emma 10, 25  
  
 Cappelen, Herman 10, 13–15, 158, 159 n.33, 203, 214, 220 n.12, 221–4, 236  
 Carnap, Rudolph 2  
 Carston, Robyn 9, 22, 35–6, 152 n.27, 158, 163, 168, 169 n.50, 172–6, 178, 183, 191–2, 195–6, 199, 200, 203–4, 228, 232  
 Chierchia, G. 53 n.30, 107 n.45, 120 n.8,  
 Chomsky, Noam 8, 40, 60 n.41, 76 n.9, 231  
 Cohen, L. J. 164  
 collective-distributive readings *see ambiguity*  
 comparison class 11–12, 26, 54–5, 83–4, 112, 122–4, 192–3, 203, 210, 216, 218, 220 n.12, 253–4, 256  
 compositionality 6, 16, 20, 25, 34, 50, 97, 105–6, 119, 127 n.11, 143–8, 155–6, 159–160, 202–3  
 Conceptual Structure 63–4  
 conditionals 11  
     indicative 15 n.1, 22, 25, 151, 170 n.52, 172–6  
     subjunctive 170–6, 204  
 contextual parameter 38, 54, 83–4, 98, 141 n.15  
 contextualism  
     about knowledge 65, 134  
     about meaning (Recanati, thesis of) 231–7, 244–5  
 Cooper, Robin 23, 56 n.34, 92, 96 n.33, 102, 107 n.45, 115, 214 n.8  
 Crimmins, Mark 37, 47 n.20, 57, 72 n.2, 162 n.37,



- Cresswell, M. 38 n.13, 49 n.23, 59, 71 n.1, 95 n.32, 96 n.33, 98 n.37, 100, 127 n.11
- Crimmins, M. 72 n.2, 162 n.37
- Culicover, Peter 59 n.38, 62–4
- deferred reference 27, 140 n.8, 206–11, 228
- definite descriptions 6, 45 n.18, 74, 87 n.23, 107, 130–1, 187–90, 252
- demonstratives 38–9, 54, 71, 73–5, 88, 99 n.39, 100, 125, 136, 138 n.4, 145–6, 149, 150, 153, 182, 187–90
- DeRose, Keith 161
- Discourse Representation Theory 43, 63–4, 96 n.34, 97 n.35, 107 n.45, 164 n.40
- domain index, *see domain variable*
- domain variable 23, 24, 90, 93–5, 100–1, 103–4, 106–7, 112–15, 117, 118, 123, 126–7, 130, 186, 214 n.8, 223–5, 248–9, 251 n.2, 252
- ellipsis  
Sellar's theory of 184, 187–91  
syntactic 40–4, 60–2, 76 n.9, 82–3, 85 n.18, 86–9, 94–5, 141 n.11, 196, 206–7, 223 n.15, 255, 259
- Engdahl, Elisabet 53 n.30, 120–1
- Evans, Gareth 43 n.15, 107 n.45
- Farkas, Donka 114 n.3, 252, 253 n.3
- Fiengo, R. 44 n.17
- free enrichment, *see pragmatics*
- Gauker, Christopher 73 n.4, 74, 75 n.8, 141 n.15, 150 n.24
- gerunds 164–5
- Gibbard, Allan 15 n.1
- Graff, Delia 59, 117–18, 119 n.6, 130, 252, 254
- Grice, Paul 3–5, 32–3, 75, 85, 90, 138, 151, 163–4, 176, 183, 193, 197, 199, 228, 231–2
- Groenendijk, J. 97 n.35, 120
- Harman, Gilbert 30
- Hawthorne, John 11, 256
- Heck, Richard 13, 14, 135
- Heim, Irene 34 n.7, 49 n.23, 54, 84, 97, 105 n.42, 107 n.45, 114, 119, 165 n.43, 241 n.3
- Higginbotham, James 34, 44, 58 n.37, 105 n.43, 164 n.39, 246
- Hintikka, J. 34, 105 n.43
- Hirst, G. 81 n.15
- indexicality 16, 19, 32, 38–40, 47, 48, 53–5, 64–5, 71, 75 n.8, 85 n.18, 87 n.23, 88, 99 n.39, 100, 135–6, 141–3, 149–50, 152–3, 155–7, 164, 182, 187, 189–90, 213, 214 n.8, 222, 234, 237–9, 243–5
- infinitives 164–5
- Israel, D. 72 n.2, 94 n.29
- it's raining* 51–4, 218, 220, 227, 240, 242–3, 257–8
- Jackendoff, Ray 59 n.38, 62–4
- Jacobson, Pauline 215, 250–1
- Janssen, T. 34 n.6
- Johnston, Mark 129, 190 n.14
- Kamp, Hans 4, 71 n.1, 96 n.34, 97 n.35, 107 n.45, 123, 254
- Kaplan, David 4–5, 53–4, 71–2, 74, 94 n.29, 135 n.1, 138, 147, 155 n.30, 233, 245
- Kennedy, Chris 237 n.11, 256
- King, Jeff 50, 97 n.35, 107 n.45, 146, 150 n.25, 154 n.29, 155, 164 n.39, 190 n.9, 204, 206, 208 n.3, 213, 233, 234 n.1, 238, 242, 245
- Klein, Ewan 76 n.9, 254–5
- location parameter 54, 226
- Kratzer, Angelika 34 n.7, 49 n.23, 54, 84, 105 n.42, 119 n.5
- Kripke, Saul 6
- Larson, Richard 119 n.5, 124 n.10
- Lepore, Ernie 10, 13–15, 141 n.15, 146 n.21, 158, 159 n.33, 203, 214, 220 n.12, 221–2, 233–4, 236
- Levinson, Stephen 158, 163, 164, 167, 168, 171–5, 178, 203, 206, 239
- Lewis, David 49 n.23, 50, 71, 105 n.42, 131 n.16, 165 n.43, 170, 219
- logical form 27, 30–2, 34, 36–7, 39, 47, 48, 52, 53, 57–65, 94, 97–8, 100, 103–4, 106–7, 113, 119, 123–4, 126–7, 158 n.32, 182, 221–2, 225, 249, 254–5, 257–8
- Ludlow, Peter 22, 162 n.37, 259

- Marti, Luisa 204, 226, 228, 244  
 mass expressions 20–1, 24, 125–9  
 May, Robert 44 n.17, 76 n.9  
 Merchant, Jason 259–60  
 minimal semantics 10–13  
 Montague, Richard 4–5, 19, 32, 71, 104, 137, 143, 231
- Neale, Stephen 83, 97 n.35, 106, 107 n.44, 187 n.6, 189 n.8, 190, 191 n.12, 200  
 Nelson, Michael 214, 222 n.14  
 nominal expression 112, 113, 117, 119, 121–4, 145 n.19, 146, 212, 213 n.7, 248–50, 254 *see also noun, quantified noun phrase*  
 nominal restriction 24, 112, 117–31, 252–3, 256  
 non-sentential assertion 22, 39–46, 258–60  
 noun 18, 19, 24, 55–6, 100–5, 107, 113, 117–18, 124, 126, 128, 130, 186 n.4, 189, 211–12, 222–5, 248–56 *see also, nominal expression, quantified noun phrase*
- Optionality Criterion 240–3
- Pagin, Peter 218–19  
 Parsons, Terence 122, 164 n.40  
 Partee, Barbara 58 n.36, 62–3, 164 n.38, 165 n.40, 203, 214 n.8, 245  
 Perry, John 24, 37, 45 n.18, 47 n.20, 51, 72 n.2, 94 n.29, 136, 139, 140 n.7, 141, 162 n.37, 183  
 Pietroski, Paul 8  
 pragmatics 10, 25–8  
   binding 192–4  
   definition 32–3, 134–2  
   free enrichment 10, 26–7, 163, 194–7, 204–5, 214–16, 238, 243, 250  
   and nominal restriction 125–6  
   and quantifier domain restriction 80–2, 85–6, 89–95, 184–6  
   and Recanati 232–44  
   truth-conditional 35–9, 46, 50–1, 62, 135  
 pragmatic binding, *see binding, pragmatic*  
 predicate 16, 19, 20, 45 n.19, 60, 65, 94, 98, 130, 146–8, 153, 154 n.29, 242, 245  
   unpronounced 86–8  
 predication, *see predicate*  
 pronominal expression, *see pronoun*
- pronoun 4, 38–9, 40, 54, 58–61, 71–2, 106–7, 113–14, 139, 152–3, 184–6, 194–5, 198–200, 215, 220 n.12, 241, 254–5
- quantified context 92–4, 100  
 quantifier domain restriction 10, 14–15, 20, 23–4, 55–7, 61, 69–70, 74–5, 81–90, 92, 93, 95, 96, 98–102, 106, 112, 114, 117–19, 122, 124, 126, 177 n.64, 184–90, 234, 242, 248  
 quantifier domain variable, *see domain variable*  
 quantified noun phrase 13, 15, 23, 100–4, 112–15, 117–19, 122, 186, 211–12, 224, 238, 249–52 *see also nominal expression, noun*  
 questions 53 n.30, 119–21
- Recanati, François 35, 47 n.20, 86 n.20, 90 n.24, 99, 135, 158, 183, 191, 192 n.13, 196 n.16, 197 n.19, 198 n.20, 200, 203, 225–6, 231–46, 257, 258  
 Reimer, Marga 73 n.3, 74 n.6, 99 n.38  
 relational expression 57–9, 61–4, 214–15, 220, 223–5  
 relevance theory 9–10, 13, 158, 163, 193, 196, 232–3  
 Reyle, Uwe 76 n.9, 96 n.34, 107 n.45  
 Richard, Mark 147 n.22  
 Rizzi, L. 39 n.14  
 Russell, Bertrand 6, 30, 136
- Sag, Ivan 208  
 Salmon, Nathan 50, 94 n.29, 138 n.4, 154 n.29, 233–4  
 Sandu, G. 34, 105 n.43  
 Saul, Jennifer 159 n.34, 176–7  
 Schiffer, Stephen 87  
 Sellars, Wilfred 184, 186–91  
 semantic modesty 158–61, 177, 204, 232–5, 237, 239, 244  
 semantic parameter approach 84, 98–9  
 semantics-pragmatics distinction 24, 32–3, 134–42  
 semantic skepticism 157–62, 177  
 Soames, Scott 31 n.2, 99, 155, 233–6  
 speaker intention 15, 24–5, 50, 74, 139, 140, 142, 145, 150, 170, 203, 205–6, 237–8

- Sperber, Dan 9, 35, 37 n.10, 47 n.20, 74, 158,  
183, 191, 192 n.13, 193, 196, 200, 203–4,  
227–8, 232
- Stainton, Robert 35, 39, 41, 44 n.16, 46, 258
- Stalnaker, Robert 12, 33, 71, 73 n.4, 74, 78,  
143–4, 172–3, 227 n.16
- Stern, Josef 207
- Strawson, Peter 3, 188–90
- strict/sloppy readings, *see ambiguity*
- Syncretic View, *see semantic modesty*
- Szabo, Zoltan Gendler 139, 140, 160–1,  
183, 196 n.16, 206, 211–12, 222 n.15,  
248, 250
- Szabolcsi, A. 38 n.13
- Tarski, Alfred 2, 19
- tense 20, 117 n.4, 164–5, 241–2, 245
- Travis, Charles 35, 80, 192 n.13, 203, 232, 237
- unarticulated constituent 47–8, 51–64,  
88–9, 93 n.28, 162 n.37, 183–4, 192–5,  
216–20, 257–258
- Varzi, A. 50
- von Fintel, Kai 23, 26, 56 n.33, 92, 100, 212  
n.6, 214 n.8
- Wettstein, H. 74 n.6, 75 n.8, 87 n.23, 141 n.14
- Williams, E., 58 n.37,
- Williamson, Timothy 11, 99, 223, 249
- Wilson, Deirdre 9, 35, 37 n.10, 47 n.20, 74,  
158, 163, 183, 191, 192 n.13, 193, 196, 200,  
203, 204, 227–8, 232
- Yanofsky, N. 39, 41