

OXFORD COGNITIVE SCIENCE SERIES

*General Editors*

MARTIN DAVIES, JAMES HIGGINBOTHAM, JOHN O'KEEFE,  
CHRISTOPHER PEACOCKE, KIM PLUNKETT

Forthcoming in the series

*Mindreading*

Stephen Stich and Shaun Nichols

*Face and Mind: The Science of Face Perception*

Andy Young

# CONTEXT AND CONTENT

*Essays on Intentionality in  
Speech and Thought*

ROBERT C. STALNAKER

OXFORD  
UNIVERSITY PRESS

**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

Athens Auckland Bangkok Bogotá Buenos Aires Calcutta  
Cape Town Chennai Dar es Salaam Delhi Florence Hong Kong Istanbul  
Karachi Kuala Lumpur Madrid Melbourne Mexico City Mumbai  
Nairobi Paris São Paulo Singapore Taipei Tokyo Toronto Warsaw  
with associated companies in Berlin Ibadan

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

© in this edition Robert C. Stalnaker 1999

The moral rights of the author have been asserted  
Database right Oxford University Press (maker)

First published 1999

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 organisation. 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

Stalnaker, Robert.

Context and content: collected papers on intentionality in speech  
and thought / Robert C. Stalnaker.

(Oxford cognitive science series)

Includes bibliographical references and index.

1. Pragmatics. 2. Intentionality (Philosophy) I. Title.  
II. Series.

B831.5.S73 1999 121—dc21 98-49339

ISBN 0 19 823708-1

ISBN 0-19-823707-3 (Pbk.)

1 3 5 7 9 10 8 6 4 2

Typeset by Hope Services (Abingdon) Ltd.

Printed in Great Britain  
on acid-free paper by

Biddles Ltd.,

Guildford and King's Lynn

*For Tom and Joanna*

natural language, but presumably intended to model, approximately, some of the devices used in natural language. But all these theories are responding to some common themes, and are motivated by the same basic ideas. I think it will help to put the theories and the phenomena they are concerned with in perspective to look back at some of the basic ideas, and to try to describe the structure of discourse in a way that abstracts away from the details about the mechanisms and devices that particular languages may provide for doing what is done in a discourse. It helps, I think, to get clear about what language is for—what it is supposed to do—before explaining how it does it.

So my intention in this discussion of the representation of context is to talk about what discourses are and what goes on in them, while saying as little as possible about the grammatical apparatus that participants in them use. I want to emphasize that it is not my purpose to criticize any of the dynamic semantic theories, or to propose alternative ways of doing what they do. My concern is with the foundations. What I will do is to sketch some descriptive apparatus that helps to display some of the ways that context and content interact. The apparatus is schematic and idealized, and is to some extent familiar (see Stalnaker, 1978, 1988), but I think it is useful to keep looking back at some of the simple truisms that underlie the sophisticated semantic and pragmatic theories that have been developed.

In Section II, I will review the motivating ideas for the descriptive apparatus that I want to recommend, and give the outlines of the apparatus itself, the central notion of which is the concept of speaker presupposition. Then in the remainder of the paper, I will try to illustrate some ways in which the application of the apparatus might help to clarify and explain some phenomena about discourse. In Section III, I will discuss the phenomenon of accommodation, arguing that the phenomenon can be explained in terms of the structure of speaker presupposition, without invoking any special rules or mechanisms of accommodation. In Section IV, I will respond to an argument given by Hans Kamp that the kind of representation of context I am recommending is not rich enough to allow for an explanation of certain linguistic phenomena. Finally, in Section V, I will consider how the framework might help to clarify some issues concerning pronouns with indefinite antecedents.

## II. *What Is a Discourse Context?*

I will start with two simple ideas that motivate the kind of representation of context that I will sketch. First, speech is action, and speech acts should

## 5

# On the Representation of Context

## I. *Introduction*

It has for a long time been widely recognized that no satisfactory semantic theory—theory of the relation between linguistic expressions and what they express—can ignore the role of the contexts in which expressions are used and interpreted. Discourse is a dynamic interactive process in which speech acts affect the situations in which they take place, and in which the situation affects the way the speech acts are understood. More than forty years ago philosophers from the British ordinary language tradition most influentially John Austin, Peter Strawson and Paul Grice—criticized traditional semantic accounts that ignored the context of the use of language and the fact that speech was a kind of action. More recently, a lot of high-powered formal machinery has been developed and deployed by linguists and logicians for the study of contexts and the discourses that take place in them—machinery that draws both on the ideas of the anti-formalist Grice–Austin tradition and on the resources of the kind of formal semantics that those philosophers were reacting to. There is situation semantics (Barwise and Perry, 1983), discourse representation theory and file change semantics (Kamp and Reyle, 1993; Heim, 1982), update semantics, dynamic predicate logic, dynamic Montague grammar (Groenendijk and Stokhof, 1990, 1991), among others. Different frameworks use different devices and apparatus, and use the apparatus in different ways to explain the phenomena. In some cases, such as discourse representation theory, the theory is presented as a semantics for English, or at least for some fragment of English, which provides an algorithm matching sequences of sentences with structures that interpret them. Dynamic predicative logic, on the other hand, is only indirectly relevant to any natural language. It defines an artificial language with new kinds of dynamic variable binding operations, obviously different from anything in

This paper grew out of material prepared for a seminar at MIT given jointly with Irene Heim. I owe much to discussion with her, and to her writings. I would also like to thank Kai von Fintel, Jason Stanley, Zoltan Szabo and several anonymous readers who commented on an earlier version of this paper.

be understood in terms of the way they are intended to affect the situation in which they are performed. Second, speech acts are context-dependent: their contents (and so the way they are intended to affect the situation) depend not only on the syntactic and semantic properties of the types of expressions used, but also on facts about the situation in which the expressions are used. These two simple observations point to two different roles that context plays: it is both the object on which speech acts act and the source of the information relative to which speech acts are interpreted. To explain the interaction of the two roles, we need a single conception of context that will represent the information about the situation that is relevant both to the role of context in determining content and to explaining how the content determined then acts on the situation. The point of a discourse—at least one central kind of discourse—is the exchange of information. Participants in a conversation begin with certain information in common, or presumed to be in common, and it is that body of information that the speech acts they perform are designed to influence. The content of an assertion will be a piece of information, and if the assertion is successful, then that information will become part of the body of information that provides the context for the subsequent discourse.

So both of the roles that contexts play require that they include a body of information: context-dependence means dependence on certain facts, but the facts must be available, or presumed to be available, to the participants in the conversation. So I propose to identify a context (at a particular point in a discourse) with the body of information that is presumed, at that point, to be common to the participants in the discourse.

Some of the common information that defines a context will be information about what the participants in the conversation are talking about: if the subject matter is the current American political scene, then speakers will take lots of familiar facts about that scene for granted (for example, that Bill Clinton is the President, and that he is running for reelection), assuming that everyone already knows them, and already knows that everyone already knows them. But the information presumed to be common will also include facts about the discourse that is currently taking place, since when one is engaged in a conversation, one can normally take for granted that speakers and hearers are aware that the conversation is taking place, that speakers are saying what they are saying in the way they are saying it. And one can take for granted that everyone is taking these things for granted. The interaction of the information of these two different kinds is a theme of much of the current work in discourse semantics and pragmatics, and it will be the central focus of my discussion here.

We can represent the information that defines the context in which a speech act takes place with a set of possible situations or possible worlds —

the situations that are compatible with the information. This set, which I have called the *context set* (Stalnaker, 1974) will include all the situations among which the speakers intend to distinguish with their speech acts. The presumed common information—what is presupposed in the context—is what all these possible worlds have in common. An assertion can then be understood as a proposal to alter the context by adding the information that is the content of the assertion to the body of information that defines the context, or equivalently, by eliminating from the context set—the set of possible worlds available for speakers to distinguish between—those possible worlds in which the proposition expressed in the assertion is false.

The world, as Wittgenstein said, is all that is the case, and a possible world is all that would be the case if that world were actual. Every proposition, relevant or not, that is taken for granted by the participants in a conversation will be true in all of the possible worlds that define the context. Since speakers in the actual world take for granted when they are talking that they are talking, they will be talking not only in the actual world, but also in each of the possible worlds that define their context.<sup>1</sup> Within each possible world in the context set, a discourse is taking place, and it has a context represented by its own context set. If, when talking about American politics, I take for granted that Bill Clinton is the President, then it will be true in each possible world in the relevant context, not only that Bill Clinton is the President, but also that I am, in that world, taking for granted that he is. Facts about what is presupposed in a context are not only facts about the actual world in which the discourse is taking place (which may or may not be a member of the context set), but also facts about the worlds that define the context. There is no problematic circularity here, but just the familiar relational structure that one finds in modal semantics, and semantics for knowledge and belief operators. The structure of speaker presupposition can be represented by a Kripke semantics in which the accessibility relation is *serial*, *transitive*, and *Euclidean*,<sup>2</sup> but not necessarily *reflexive*. The requirement that the relation

<sup>1</sup> As Zoltan Szabo has pointed out to me, just what a speaker, or a writer, presupposes about the situation in which she is speaking or writing, will vary from case to case. While in an ordinary face-to-face conversation, the speaker will presuppose that her addressee hears her as she speaks, and knows the time of utterance, this won't be true for answering machine messages. To interpret a particular token of an answering machine message like "I can't come to the phone now," one needs a context in which the addressee, but not the speaker, is presumed to know the facts necessary to determine the content of the message.

<sup>2</sup> A binary relation  $R$  is serial if:  
 $(x)(\exists y)(Rxy)$   
 and it is Euclidean if:  
 $(x)(y)(z)((Rxy \wedge Raz) \rightarrow Ryz)$ .

be transitive and Euclidean reflects the assumption that speaker presupposition is transparent: speakers know what they are presupposing, so they presuppose that they are presupposing that  $P$  if they are, and that they are not if they are not. The requirement that the relation is serial is simply the assumption that the context set is always nonempty—there is always at least one possibility compatible with what is being presupposed.

The nonreflexivity of the relation reflects the important fact that the actual world in which a discourse takes place need not be compatible with the context of that discourse, which is just to say that some things presupposed by a speaker may be in fact false. This may happen for a number of different reasons. A speaker might be presupposing something false simply because he has a false belief, either about the subject matter of the discourse, or about the discourse itself. He might presuppose something false in order to deceive, or because of some mutually recognized pretense. Sometimes the most effective way to communicate something true is to presuppose something false. For example, if you are presupposing something false but irrelevant, I may presuppose it as well, just to facilitate communication. (You refer to Mary's partner as "her husband," when I know that they are not married. But I might refer to him in the same way just to avoid diverting the discussion.)

I have said that we should identify a context, or context set, with the set of possible situations that are compatible with what is presupposed, or taken to be common ground, by the participants in the discourse. But it is individuals who presume that something is or is not common information, and different participants in the same conversation may not in fact agree in what they presuppose. But even if actual presuppositions of speaker and addressee diverge, it is part of the idea of speaker presupposition that it is presupposed that the presuppositions of the different participants in the discourse are the same. To presuppose something is to presuppose that it is common ground, so if what is presupposed is not in fact common ground, then something false is being presupposed. So in those possible worlds compatible with what anyone is presupposing, all the participants (at least everyone that the speaker presupposes to be a participant) will be presupposing the same things. This fact about the concept of presupposition will be reflected in a constraint on the accessibility relations in the modal semantics for presupposition, in this case a constraint on the relation between the different accessibility relations for the different speakers. The following is the relevant condition: for any two participants  $i$  and  $j$  in the discourse, and any possible worlds  $x$ ,  $y$  and  $z$ , if  $R_{i,x}$ , then  $R_{j,z}$  iff  $R_{j,y}$ .

It would, of course, be unrealistic to assume that the participants in a discourse are always in fact presupposing exactly the same things, and the

constraint I have proposed does not assume this. This presupposition, like any presupposition, may be false in the actual situation. We need not even assume that any participant *believes* that this presupposition holds, since this presupposition, like any presupposition, may involve the pretense that something is common ground when it is not. But the constraint is important, since it will help to explain some of the ways in which contexts change in response to events that take place in a discourse. Because it is presumed that shared information is shared, the presuppositions of different speakers will tend to adjust to an equilibrium position in which they are the same.

Discourse contexts, I have been suggesting, can be represented by the set of possible situations compatible with the information that is presumed, by the speaker, to be common ground, or information that is shared and recognized to be shared by all the relevant participants. This is a representation from the point of view of one of the participants in the context of what is common to all, and it is a representation of a context at a particular moment. This body of information is the information that is available for the interpretation of what is said in the context. If certain information is necessary to determine the content of some speech act, then appropriate speech requires that the information be presumed to be shared information at the time at which that speech act is to be interpreted. But exactly what time is that? The context—what is presupposed in it—is constantly changing as things are said. The point of a speech act—an assertion, for example—is to change the context, and since the way the speech act is supposed to change the context depends on its content, interpretation must be done in the prior context—the context as it is before the assertion is accepted, and its content added to what is presupposed. But the prior context cannot be the context as it was before the speaker began to speak. Suppose Phoebe says "I saw an interesting movie last night." To determine the content of her remark, one needs to know who is speaking, and so Phoebe, if she is speaking appropriately, must be presuming that the information that she is speaking is available to her audience—that it is shared information. But she need not presume that this information was available before she began to speak. The prior context that is relevant to the interpretation of a speech act is the context as it is changed by the fact that the speech act was made, but prior to the acceptance or rejection of the speech act. A successful statement will thus change the context in two different ways that need to be distinguished. First, the fact that the statement was made is information that is added to the context simply as a result of the fact that it is a manifestly observable event that it was made. Second (assuming the statement is not rejected), the content of the assertion will be added to the context. One might imagine a formal language

game in which the two kinds of change take place successively: first, the addressee takes in and adjusts to the fact that a particular proposition has been asserted—that it has been proposed that a certain piece of information be mutually accepted; second, the addressee accepts or rejects the proposal, either adding the content of the assertion to the contextual information, or leaving the context as it was after the first change. But there is no reason why the two changes need take place at different times. In real conversations, particularly conversations that are exchanges of information by people whose authority is not in question, there need be no moment when the context has changed in the first way, but not yet in the second. Nevertheless, even when the two changes take place at once, the distinction between them is important, since the explanations will be different for the different kinds of changes. I will try to bring out the significance of the distinction by discussing, first, the concept of *accommodation*, second, by responding to an argument that the possible worlds representation of context is not sufficiently rich to account for some phenomena, and third, by looking at some simple examples of pronouns and indefinite reference.

### III. Accommodation

Context changes of the first kind include the recognition of unproblematic observable events, such as that a certain person uttered certain words, but they also include the recognition of any facts that the hearer can infer from the occurrence of these observable events, with the help of standing information that is part of the prior context. We add to the context, not only the information that Phoebe uttered certain sounds, but also that she uttered an English sentence, and that she is saying something to us. And if it was part of the information presupposed before she began to speak that Phoebe is a cooperative and competent participant in our conversational enterprise, and if her utterance does not give us reason to give up this presupposition, then we can also infer that she is making whatever presuppositions are required to make her utterance intelligible and appropriate. The phenomenon is familiar: Phoebe says

(1) I can't come to the meeting—I have to pick up my cat at the veterinarian.

Her utterance is appropriate, let us suppose, only if it is presupposed, and so part of the presumed shared information, that Phoebe owns a cat. But suppose Phoebe had no reason to believe that her interlocutor knew about her cat before she made her statement. That does not matter, since Phoebe

can reasonably assume that her interlocutor can infer that she has a cat, and that she is presupposing that she has a cat, from the prior presupposition that she will speak appropriately, and from the fact that her statement would be appropriate only if she were making that presupposition. This phenomenon, which David Lewis has labeled "accommodation," is just a case of inference from the observed fact that the utterance event takes place. Lewis, in setting out a dynamic model of a conversational game, described a distinctive kind of rule, a "rule of accommodation," that is characteristic of linguistic practices (Lewis, 1979). Such rules, he says, ensure that presuppositions "straightway come into existence" when they are required to render what someone says acceptable. Lewis distinguished context changes by accommodation from the simpler kind of change that takes place when "something conspicuous happens at the scene of a conversation, and straightway it is presupposed that it happened," but I do not think there is a distinction to be made here, or a reason to hypothesize a special rule or mechanism of accommodation. An utterance event is something conspicuous that happens at the scene of a conversation, and the presupposition that such an event occurred is the source of any accommodation.

To make this point clearer, let me sketch in a little more detail how the derivation of the accommodated presupposition might go. First, suppose that it is a prior standing presumption that Phoebe is a competent and cooperative speaker, so it is presupposed that her speech acts will be appropriate. Second, suppose it is also a prior standing presupposition that the definite description, "my cat" in the expression "I have to pick up my cat" presupposes that the speaker has a cat, in the sense that the use of this expression is appropriate only in a context in which the speaker is making this presupposition—presuming that it is shared information that the speaker has a cat. This, we can assume, is something that competent speakers of English know, and it is a standing presupposition in an ordinary conversation taking place between English speakers that they are all competent speakers. Now if these standing presuppositions remain in force when we add the new information that Phoebe uttered (1), then it will follow that Phoebe is presupposing that she has a cat. But since what is presupposed is what is presumed to be shared information, this implies that Phoebe is presupposing that her addressee is presupposing that she has a cat. (Not that her addressee was presupposing it before she began to speak, but that he is presupposing it after having taken in the fact that she made her statement.) So since it follows from the addressee's standing presuppositions together with manifestly observable facts that Phoebe is making this presupposition, he will presuppose that Phoebe is making it, and since it is presupposed that presuppositions are shared information, he will accommodate by presupposing it himself.

The last move in this derivation requires that the addressee be willing to suppose that if Phoebe is presupposing that she has a cat, then it is true that she has a cat, and so the appropriateness of Phoebe's presupposition will require that she be presupposing that the addressee will be willing to suppose this. If this assumption is not correct—if something controversial has been presupposed (either as a result of the speaker's misperception, or in a deliberate attempt to sneak something by the addressee)—then the context will become defective, and some backtracking and repair will be required. So our derivation helps to explain a familiar feature of accommodation: that (to quote Irene Heim), "assumptions to be accommodated are supposed to be uncontroversial and unsurprising. One may explicitly assert controversial and surprising things (in fact one should) but to expect one's audience to accept them by way of accommodation is not good conversational practice" (Heim, 1992: 212).

#### IV. *On an Alleged Problem with This Representation of Context*

Hans Kamp has argued that this framework for a representation of context is not rich enough to allow for some distinctions that are needed to account for some linguistic phenomena. I don't think his argument is correct, and I think the explanation of why it is not correct helps to bring out the significance of the two contrasting ways in which speech acts change the context. More generally, I think it will help to clarify the contrast between the descriptive apparatus that I am sketching and a theory of the mechanisms of a particular language, such as discourse representation theory. I will argue that while the example on which Kamp bases his argument does point to a fact about the language that needs explanation, it does not point to an inadequacy in the representation of context I have proposed, which is designed not to provide a solution to that problem, but to provide a framework in which the problem, and alternative responses to it, can be clearly formulated.

Kamp's argument (Kamp, 1988) begins with a pair of examples due to Barbara Partee. Consider these two contrasting sentence pairs:

- (2) a. Exactly one of the ten balls is not in the bag.  
b. It is under the sofa.
- (3) a. Exactly nine of the ten balls are in the bag.  
b. It is under the sofa.

Kamp asks us to assume that each discourse takes place in the same initial context. In each case, the first assertion changes the context, and the second assertion is made in the changed context. He then notes that the first

statements of each pair, (2a) and (3a), are truth-conditionally equivalent—true in exactly the same possible worlds. So he concludes that if contexts are simply identified with sets of possible worlds, and if assertions simply add the truth-conditional content to the context, then the posterior context that results from the assertion of (2a) will be the same as the one that results from the assertion of (3a). But on the other hand, the contexts in which the second statement takes place must be different in the two cases, since (2b) is the same as (3b), yet they cannot receive the same interpretation. The "it" in (2b) can, and most naturally does, refer to the ball that is not in the bag, while the "it" in (3b) cannot refer to that ball. "We must conclude," Kamp says, "that no difference can be predicted if contexts are identified with sets of possible worlds. Therefore, a theory of meaning and context dependent interpretation of English must, if it is to handle such examples successfully, adopt a representation of contexts that goes beyond what sets of possible worlds are able to reveal" (Kamp, 1988: 158).

Now it is right that the abstract framework does not *predict* this difference, since it says nothing about the way pronouns work. This representation of context is not intended as a part of a theory of meaning for English, or any other language, natural or artificial, but as a piece of apparatus with which the phenomena might be described and explanatory hypotheses stated. But does the framework permit a distinction between the contexts for (2b) and (3b), given that the prior contexts for (2a) and (3a) are the same? Kamp's claim is that the representation of context by a set of possible worlds is not sufficiently rich to allow the phenomena to be explained, but this is not correct. One cannot conclude from the fact that the prior context (the context as it was before the speaker began to speak) is the same for (2a) and (3a), and the fact that the truth-conditional content of (2a) and (3a) are the same, that the resulting contexts (the contexts for (2b) and (3b)) must also be the same. The assumption that one can draw this conclusion ignores the first way in which a speech act changes the context. Since it is a manifestly observable fact that, in each case, a certain sentence was uttered, this fact, together with any additional information that follows from that fact, conjoined with standing background information about linguistic and speech conventions, is available to distinguish the two posterior contexts, the contexts relative to which (2b) and (3b) are interpreted.

A pronoun such as "it" presumably requires a context in which a certain individual of the appropriate kind is uniquely salient, or in some way available for reference. The abstract framework, by itself, says nothing about which particular facts make salient, or available for pronominal reference, an individual who inhabits the domains of the possible worlds that define the context. Among the facts that might be relevant are facts about the

words used to express certain propositions in the discourse, since these are facts that are available to the participants in the discourse, and so are reflected in each of the possible worlds that are members of the relevant context set. So long as we make the minimal assumption that information can be relevant for the determination of content only if it is presupposed by the speaker to be information that will be available to the addressee, then we can be sure that the set of possible worlds that defines what is presupposed will suffice to represent a context.

I want to emphasize that I am not suggesting that Partee's example does not pose a real problem: the problem of explaining just what the relevant difference is between (2a) and (3a), the difference that is responsible for the fact that (2a) makes a certain individual available for pronominal reference, while the truth-conditionally equivalent (3a) does not. Discourse representation theory and dynamic predicate logic each have something to say about what the mechanisms that explain the difference might be. But whatever the explanation, it will appeal to facts about English, or about the practice of speaking English, that can be presumed to be available to competent English speakers, and so they will be facts that distinguish the possible worlds that define the different contexts.

I also want to emphasize that in giving this response to Kamp, I am not proposing to add anything new to the possible worlds that define a context, and am not proposing any alteration in the straightforward account of the content of sentences such as (2a) and (3a). I am simply observing something obvious: that when a speech act takes place, it is a manifest fact, and so at that point presumed shared information, that it takes place. It is simply an obvious fact that an assertion changes the context, not only by adding the information that is the content of the speech act (the same information in the case of (2a) and (3a)), but also by manifesting the fact (different in the utterances (2a) and (3a)) that the event itself takes place.

#### V. *Pronouns and Indefinite Reference*

To try to get clearer about the way a context might make individuals available for reference, I will look at some simple examples of pronouns with indefinite antecedents. Start with the following unproblematic opening move in a conversation.

- (4) I met an interesting woman at the dinner last night. She was a member of Clinton's cabinet.

One fact about the context that results from the acceptance of the first of these two statements (a fact about the subject matter of the discourse) is

that in all of the possible worlds in the context set, the speaker met an interesting woman. No particular woman was specified, so it need not be the same interesting woman that the speaker met in each possible world in the context, and nothing rules out the possibility that the speaker met several interesting women in some of the possible worlds compatible with the information. A second fact (a fact about the discourse) that is presupposed after the first statement has been made is that a woman that the speaker met at the party last night is uniquely available to be the referent of the pronoun "she" in the conversation in question. At least that must be true if the second statement is to discriminate in a determinate way between the possible situations that define the context. But if there are possible worlds compatible with the context in which the speaker met two interesting women at the dinner, what makes one or the other of them the referent of "she" in that world?

Presumably, in this particular example, the speaker has a particular woman in mind,<sup>3</sup> but since the audience knows nothing about the woman at the dinner except what the speaker tells them, and since the speaker has so far given them only general information, it cannot be a fact about the context—about what is presumed to be common knowledge—that any particular woman is salient, or available to be referred to in the subsequent conversation. We can, however, say that it is a fact about the individual possible worlds that are the members of the set that defines the context that a certain person is available to be the denotation, in that world, of a pronoun. This will be a fact about the conversation, and not about the subject matter. There might be two possible worlds compatible with the context that are indistinguishable with respect to the events that happened at the dinner, but different with respect to which of the two interesting women is the potential denotation of the pronoun "she." Suppose it is compatible with the context, after the first of the two statements, that the speaker met an interesting woman who was a member of Clinton's cabinet, and also an interesting woman who was not. Then there will be two possibilities corresponding to this situation: one in which the first interesting woman is marked as the potential referent, and one in which the second is. When the speaker goes on to make the second statement, its content eliminates the second of these possibilities. While this difference between possibilities does not concern facts about the subject matter, it is

<sup>3</sup> Just what it means to say that one has a particular person or thing in mind is, of course, notoriously difficult to pin down. It is obviously itself a highly context-dependent notion. I don't take the phrase to imply that one has a very intimate knowledge of the individual, or that one has the capacity to identify the individual, by description, or by observation. Even if our speaker would now be unable to distinguish the woman he is talking about from another woman he met at the party, it might still be true that there is a particular woman he is talking about, presumably the causal source of his beliefs about her.



a factual difference: the fact that distinguishes the two possible worlds is a fact about the discourse which is taking place, not only in the actual world, but also in each of the possible worlds compatible with the context. Specifically, it is a fact about the interpretation of certain potential continuations of the discourse.

The recognition of the possibility that the speaker met two interesting women might be quite explicit in the context. Suppose the speaker had begun his conversation this way:

- (5) I met two interesting women at the dinner last night. Let me tell you about one of them.

Here the second speech act conveys no information about the events at the dinner. All it does is to make available one of the two women for reference by a pronoun. Which one? One of them in some possible worlds compatible with the context, and the other in others. But what sort of fact about the world is it that determines which of the women it is? It is clear enough how the facts determine whether what is presupposed or said about the events at the dinner is correct or not. The speaker either did or did not actually meet an interesting woman who was a cabinet officer. But if the speaker really did meet two interesting women, in a certain possible world, how do the facts about that world determine which of them is the denotation of the speaker's "she"? At least in this particular example, it seems reasonable to say that this is up to the speaker. Suppose the actual facts are these: one of the two interesting women that our gullible and politically unaware speaker met was a professor of philosophy who was introduced to him by a practical joker as the Secretary of Defense. The other was in fact the Secretary of Health and Human Services, but the speaker was ignorant of this fact—he thought she was a talk show host. So the speaker correctly believes that he met exactly one interesting woman who was a cabinet officer, but he is mistaken about which of two women it was. The facts about which the speaker is mistaken will not affect what he presupposes, or how what he says changes the context, but it will affect how his statements are assessed, and how they might be corrected. A participant in the conversation who was aware of the relevant facts might correctly respond to our speaker's second statement by saying

- (6) She wasn't a member of Clinton's cabinet—that was a joke.

So it seems clear that if the facts were as described, the speaker's second statement would be false. In contrast, if he had instead told his story in a single statement, saying

- (7) I met an interesting woman who was a member of Clinton's cabinet at the dinner last night

then he would have spoken truly, even if only by a lucky accident. But as he did tell it, it seems to me that he said something false.<sup>4</sup>

In early work on formal pragmatics by Richard Montague and some of his colleagues, contexts were identified with *indices*, n-tuples that included all the features of the situation on which the extensions of expressions might depend. So an index would contain a speaker, an addressee, a time, a place, a possible world, and other things such as "indicated object" coordinates to determine the referents of demonstrative pronouns. Meanings were identified with functions from such indices to extensions (see Montague (1974), Lewis (1970, 1981)). David Kaplan's work on demonstratives suggested that meaning takes context to content, and content takes circumstances to extension. But contexts were still represented by n-tuples of elements (speaker, time, indicated objects, etc.)—everything other than meaning that is needed to determine content. Context, on this account, was constituted by those elements required to determine what was being said with a sentence, but nothing was explicitly said in the theory about the epistemic status of such a context. If, however, we assume that the speaker presupposes that the addressee will know what is being said, then we will be assuming that the addressee knows what the relevant index is, which implies that the "indicated objects" must be salient, or that the speaker must be assuming that everyone knows which individual a pronoun denotes. But all that is really required for content to be determined, and for it to be known in the context how it is determined, is for there to be a unique available "indicated object" in each possible situation compatible with the context. What the context must determine, for the interpretation of a pronoun, is a function from worlds in the context set to individuals.

So to return to our example of an opening conversational gambit, the idea is that the first statement of (4) changed the context in two ways: by adding the purely existential information, and by making, in each possible world compatible with the resulting context, a certain individual available for future reference by "she." If the statement is rejected, then the change induced by the existential information will not take place, but in some cases, this will not prevent the other change—the making available for

<sup>4</sup> Kai von Fintel pointed out to me that (6)—the correction—would be an appropriate response to (7) as well as to (4), and so it is not clear that (7) is true. I agree that the correction would be appropriate in either case, and true in both cases, but while I think the truth of (6) shows that the second statement of (4) is false, I don't think it shows that (7) is false. When the speaker of (6) says "she wasn't a member of Clinton's cabinet," he directly contradicts the second statement of (4). "she was a member of Clinton's cabinet." So if (6) is true, the second statement of (4) must be false. But while (6) gives a reason to reject (7), it doesn't follow that it contradicts it. Compare the Strawson correction example discussed below.

future reference of a certain individual. Consider a famous example of Peter Strawson's (1952: 187):

- (8) X: A man jumped off a bridge.  
Y: He didn't jump, he was pushed.

In this dialogue, X succeeds in making a certain man available for reference (by Y) in the subsequent context, even though the statement that accomplished this was rejected by Y. It was the fact that the statement was made, and not the fact that it was accepted, that did the job.

It is the speaker introducing the individual into the context, and not necessarily the individual using the pronoun, whose intentions are relevant to determining the referent of the pronoun. In Strawson's correction dialogue, Y's "he" refers to the person X had in mind. The dialogue might take the following turn: After Y says "he didn't jump, he was pushed," X corrects the correction:

- (9) No, I didn't mean that man. I agree he was pushed, but the one I was talking about really did jump.

Even though Y had a particular person in mind who he believed to be the referent of his "he," it is the person X had in mind to whom he refers, and that is why Y can be corrected. (But notice that even though Y's correction is rejected, he succeeded in changing the context by making a different man available for reference by X.)

It will of course not always be true that a speaker making an existential statement has a particular individual in mind. Suppose Z says:

- (10) I predict that a woman will be nominated for President in 2000. Furthermore, I predict that she will win.

Suppose (to Z's surprise and delight) both parties nominate women in 2000. Naturally, one wins and one loses. Was Z's statement true or false? It would seem churlish to say that her prediction was anything but completely correct, but it is not clear that our semantic intuitions, strictly speaking, support this judgment. There does not, in this case, seem to be any fact that makes one or the other nominee the referent of the speaker's "she." Perhaps, in this case, the speaker was presupposing that at most one woman would be nominated, if only because the alternative did not occur to her. If that is true, then the second statement will have a determinate truth-value in each possible world in the context set, and so will change the context in a determinate way, but it will not have a determinate truth-value in the actual world.

Linguists sometimes distinguish specific from nonspecific indefinite noun phrases, corresponding roughly to the case where the speaker has

someone or something in mind, and cases where he does not. The problem has always been to understand what difference this distinction makes for the truth-conditions of statements containing the indefinite expression, and how facts about the speaker's state of mind that were unavailable to the audience could be relevant to the semantics of what was being said. The account I am sketching suggests that this difference matters, not to the interpretation of the indefinite expression itself, but only to the evaluation of subsequent statements made with pronouns anaphoric to the indefinite expression.

Assertions, I have suggested, are proposals to change the context by adding the information that is their content. We might contrast assertions with a different speech act, supposition, that has a similar effect, but a different point. A supposition—what one does, for example, with the antecedent of a conditional—is different from an assertion in two obvious ways: first, the intention is to add the content expressed in the supposition to the context only temporarily; second, an act of supposition does not represent its content as something the speaker believes, and is not subject to criticism on the ground that it is false. But since it alters the context, at least while it is in effect, in the same way as a successful assertion, one should expect to find parallels in the behavior of the two speech acts, and one does. One also finds some contrasts, which I think can be explained in terms of the two obvious differences between the speech acts. Suppositions change the context in the same two ways that assertions do, but in the case of supposition, while the change induced by the content is temporary, the preparatory change—the change induced by the fact that the supposition was made—is not. We can, for example, make a particular individual salient by referring to him in a supposition, and he may remain salient even after we are no longer making the supposition. So I can say, for example,

- (11) If Colin Powell runs for President, I will be surprised. I won't, however, be surprised if he is nominated for Vice President.

Even when an individual is introduced into the context by an indefinite reference in a supposition, an individual may sometimes be made available beyond the scope of the supposition. Consider:

- (12) If a man with a Scandinavian accent calls, tell him about the leak in the drain pipe under the kitchen sink. He's the plumber. I left a message on his answering machine.

The situation here parallels the Strawson example of correction. Just as, given the right background information, an individual introduced in an assertion may remain available for reference even when the assertion is rejected, so in some contexts an individual introduced in a supposition

may remain available for reference after the supposition is no longer in force.

In the case of suppositions made for the purpose of contingency planning or hedged predictions, the content of the supposition is simply added, temporarily, to the context; no presuppositions are subtracted, or suspended. But with counterfactual suppositions, one requires a new temporary context set, disjoint from the primary one. One crucial difference between an ordinary context and a counterfactual or derived context—a context brought into play either by a counterfactual supposition or some other embedded expression of a proposition, such as a belief attribution—is that in the possible world compatible with a derived context the discourse itself need not be taking place. If X says to Y, "Suppose I had never met you," the possible worlds compatible with what is being supposed are not worlds in which their conversation is taking place. So derived contexts do not change in the same two ways in the course of the conversation. If an individual in the possible worlds in a derived context is uniquely available for reference, this will have to be a fact about the conversation taking place in the possible worlds in the primary context. In interpreting speech acts with embedded sentences, the interplay between information represented by primary and derived context sets can help to clarify phenomena concerning intentional identity and *de re* belief attributions. But that is a discussion for another time.

## VI. Conclusion

My main aim in this paper has been to redirect attention from the details of the devices and mechanisms of particular natural languages to the general structure of the practices in which natural languages are used to serve the ends of speakers and listeners. The hope is that if we get clearer about the structure and purposes of discourse, we can better distinguish the idiosyncratic features of particular conventional devices from more general features of the practice that follow from assumptions about what people engaged in it are trying to do. With a framework for representing the general features of discourse, we should be able to describe in a perspicuous manner the problems that alternative semantic theories are addressing, and the alternative hypotheses about how different languages and discourse practices work. And perhaps if we are clearer about the general structure of discourse this will help us defend simpler semantic analyses. The general strategy is like the one proposed by Paul Grice and spelled out in the beginning of his William James lectures about thirty years ago. Grice's more specific project at that point was to reconcile traditional

semantic analyses of logical expressions—those offered by truth-functional logic—with the diversity of facts about the uses of such expressions by invoking some simple maxims of conversation, maxims that assumed little more than that conversation was a rational and cooperative activity whose primary purpose was the exchange of information. The semantic theories for natural language available today are considerably more complex and sophisticated than those that Grice was concerned to defend, and a lot more has been said since then—much of it inspired by Grice's ideas—about the general structure of conversation. But I think the general Gricean strategy of trying to reduce the burden on semantics by explaining as much of the phenomena as possible in terms of truisms about conversation as a rational activity remains as fruitful and promising as it was when he first proposed it.