



OXFORD



# Occasion-Sensitivity

*Selected Essays*

CHARLES TRAVIS

## OCCASION-SENSITIVITY

*This page intentionally left blank*

# 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

© Charles Travis 2008

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

First published 2008

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 the same condition on any acquirer

British Library Cataloguing in Publication Data  
Data available

Library of Congress Cataloguing 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 978-0-19-923033-4

1 3 5 7 9 10 8 6 4 2

# Occasion-Sensitivity

*Selected Essays*

CHARLES TRAVIS

**OXFORD**  
UNIVERSITY PRESS

# *Contents*

<i>Acknowledgements</i>	vii
Introduction	1
Part I. Occasion-Sensitivity	
1. On What Is Strictly Speaking True	19
2. Annals of Analysis	65
3. Meaning's Role in Truth	94
4. Pragmatics	109
5. Sublunary Intuitionism	130
6. Insensitive Semantics	150
7. Aristotle's Condition	161
Part II. Applications	
8. Are Belief Ascriptions Opaque?	185
9. Vagueness, Observation, and Sorites	206
10. Attitudes As States	227
11. On Concepts of Objects	253
12. On Constraints of Generality	271
13. A Sense of Occasion	290
<i>References</i>	316
<i>Index</i>	319

*This page intentionally left blank*

## *Acknowledgements*

The following material has been previously published, and is reprinted here with kind permission.

'On What Is Strictly Speaking True', *The Canadian Journal of Philosophy*, v. 15, n. 2, June 1985, pp. 187–229.

'Annals of Analysis', (Critical notice of H. P. Grice, *Studies in the Way of Words* (Cambridge, Massachusetts, Harvard University Press), 1998) *Mind*, v. 100, n. 2, April 1991, pp. 237–264.

'Meaning's Role in Truth', *Mind*, v., n. 3, 105, July, 1996, pp. 451–466.

'Pragmatics', *A Companion to the Philosophy of Language*, Bob Hale and Crispin Wright eds., Basil Blackwell, Oxford, 1997, pp. 87–107.

'Sublunary Intuitionism', *Grazer Philosophische Studien*, vol. 55, 1998, pp. 169–194.

'Insensitive Semantics', (Critical notice of H. Cappellen and E. Lepore, *Insensitive Semantics* (Oxford: Basil Blackwell), 2005) *Mind and Language*, v. 21, n. 1, February 2006, pp. 39–49.

'Aristotle's Condition', to appear in *Timothy Williamson and his Critics*, P. Greenough and D. Pritchard, eds. (Oxford: Oxford University Press), (forthcoming).

'Are Belief Ascriptions Opaque?', *Proceedings of the Aristotelian Society*, vol. 85, 1984/5, pp. 73–100.

'Vagueness, Observation, and Sorites', *Mind*, v. 93, n. 3, 1985, pp. 345–366.

'Attitudes As States', *Proceedings of the Fourteenth Wittgenstein Symposium: Centenary Celebration*, R. Haller and J. Brandl, eds., (Vienna: Hölder-Pichler-Tempsky), 1990, pp. 265–284.

'On Concepts of Objects', *Karlovy Vary Studies in Reference and Meaning*, Hill and Kot'atko, eds., (Prague: Filosofia), 1995, pp. 365–389.

'On Constraints of Generality', *Proceedings of the Aristotelian Society*, v. 94, 1994, pp. 165–188.

'A Sense of Occasion', *Philosophical Quarterly*, v. 55 n 219, April 2005, pp. 286–314.



*This page intentionally left blank*

## *Introduction*

When would what we say be true? That is the central question to which occasion-sensitivity is an answer. The first part of this collection consists of contributions to working out that answer.

If I am right, occasion-sensitivity is everywhere. In quotidian affairs it rarely attracts, or needs, attention. It is as discreet and unassuming as, say, the principles of pronoun binding. It is another matter in philosophy. Not, it seems, that it clamours for attention. More that we neglect it at our peril. For it matters very much to the questions philosophy has raised, and to what questions it should raise—to what there is to ask about. The second part of this collection consists of applications of the idea of occasion-sensitivity to a range of specific, and venerable, problems. Some of these are parts of ongoing endeavours, others I wish were.

When *would* what we say be true? Part of an answer might be: only when it aims at truth. Perhaps not even that is right. But it gestures at how, in these essays, I have (for the most part) narrowed the field of inquiry. Beyond that, it is all too easy to think that one has answers to the question in advance of looking. We say the things we do (for the most part) in a language. The language we use (the words we utter) speaks of various ways for things to be; we use it to speak of various things as those ways. What we thus do, one thought is, determines when what we say in those words would be true. Which encourages the thought that if we want to say what an open sentence of, say, English means, we can do this by saying when it would be true of a sequence of things that it might jointly speak of. Or if we want to say what a closed sentence means, we can do this by saying simply when it would be true (of the null sequence).

But suppose we approach the question with an open mind. We might look to see how, in particular cases, one thing we might say *does* differ from another in when it would be true, or by what it might differ. The features which mark the differences we thus find are all ones one might need in order to identify a truth-bearer; all ones on which the truth of some of what we say (or might) depends. This may, or may not, support the view that what the words we use mean determines (modulo referents) when they would be true. I called that a view. Some see it as a truism. So posing the question makes it a substantial claim.

I found the idea of approaching language in this way in Chomsky. But Austin had it, and applied it. His first and main result is set out in *Sense and Sensibilia* as follows:

It seems to be fairly generally realized nowadays that if you just take a bunch of sentences . . . impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for (leaving out of account so-called ‘analytic’ sentences) the question of truth and falsehood does not turn only on what a sentence *is*, nor yet on what it *means*, but on, speaking very broadly, the circumstances in which it is uttered. Sentences are not *as such* either true or false. (Austin 1962*b*, pp. 110–11)

Austin proved badly mistaken about the ‘generally realized’ part. He was right about the rest. That rest is the core idea of occasion-sensitivity.

To see the depth of Austin’s point, one can best frame it as about open sentences. Take a closed sentence like ‘Sid is at home’. ‘Sid’ might be used to refer to any of many people. The verb, perhaps, refers to time of speaking. So when what *this* sentence says would be true cannot be decided just by what it means (or its parts do). That is the shallow point. It is shallow because it allows us still to think that we can identify in this closed sentence an open one, with places for reference to the various items—a person, a time—that would need to be spoken of in using it if one is to say something true or false. (The open places in this open sentence may or may not all correspond to words actually in the closed one. For example, there is no *explicit* mention of a time there. No matter for the present point.) The idea is: if we choose our open sentence well, then when *it* would be true of a *sequence* of items—each, respectively, a potential referent of what filled its corresponding open places—*is* determined by the meanings of the words. The words ‘is at home’, say, would be true of a person and a time (the thought is) just in case that person is home at that time. So the spirit, if not the letter, of the law remains.

But this is *not* Austin’s point. It is rather: take any *open* sentence you like, with any number of supposed places in it, each to be filled with reference to a given sort of thing. Take any sequence of things, each fit for reference in its corresponding place. Then what the open sentence *means* is compatible with it saying any of indefinitely many different things of *those* things, so referred to in some closing of it. For example, there is indefinite variety in the things to be said in saying someone to be home at a time. (Are you at home when your house, you in it, has just slid down the hill?)

In the English sentence ‘The sky is blue’, ‘the sky’ speaks of the sky, ‘is blue’ (roughly) of being blue. If you are relaxed enough, or, perhaps, sensitive enough, you might see ‘the sky’, on a given use of this sentence, as speaking of the sky at, or about, some contextually definite place. That place where the skies are always blue, say. Perhaps (or perhaps not) the ‘is’ in ‘is blue’ speaks of some particular time. Now we have fixed some objects, and ways for them to be, which are

referred to, or spoken of, in that sentence *per se*. This much, and no more, is what that English sentence does in meaning what it does. But, since there are many things to be said in speaking of those things, it must then be that being blue, in this case, admits of understandings: there are various things one *might*, at least sometimes not incorrectly, take being blue to be. Being blue might require of the sky that when one is in it (in a plane, say) one be enveloped in blue, or it might not require this. When one calls the sky blue, e.g., in saying, ‘The sky is blue’, one might thus speak on one or another understanding of its so being; its being blue on that understanding being what it is for things to be as one thus said. This illustrates the point at work.

When I say the sky to be blue, I speak of the sky, express the concept of being blue, and represent the sky as what that concept is a concept of. For all of which I may have said any of many things, each different from the others in when it would be true. This is one way of saying what I just said otherwise above. It means this: if there are such concepts as that of being blue, that of being a sirloin, that of being a hamster, and so on—if that is the sort of thing a concept is—then concepts admit of many applications, and are satisfied (or not) only *on* one such or another. There are many understandings of being what the concept is of being; there are, correspondingly, many understandings of satisfying the concept. (One can think of this as anti-Fregean. One can also think of it as suggesting a new way of reading Frege. But this last idea is a story for another place.)

I have presented this as what is most central to Austin’s view of language. Well, almost. One might also seek what is most central to Austin in *How To Do Things With Words*—a work that has been enlisted in some rather different projects. My starting point puts a particular gloss on that work. I read it as an extended study of the concept truth. Part of Austin’s aim in it, he said, was to ‘play old Harry’ with the fact/value distinction. Part of his way of doing that was to detail just how similar truth-bearers are to speech acts which do not aim at truth; just how similar being true is to the other successes those other speech acts aim at. This is the moral Austin draws at the end of that work. Of course speaking truth is saying things to be as they are, and none other. But when a question arises as to whether someone has spoken truth, say, in telling us there are biscuits on the sideboard, questions *may* arise as to what one wants to count as doing that. If the biscuits are months old, or dog biscuits, does that count? Circumstances of evaluating a statement may then matter to an answer.

To state something is to aim at truth. Where there is room for question as to what *ought* to count as reaching this aim—where success is not univocally decided just by the world being as it is—there is also room for a statement to adjust its aim accordingly. If mouldy biscuits *may* count as success for a statement of there being biscuits on the sideboard, then it may be built into a given statement that, for it, at least, this *will* count as success. It may be so to be understood. It may be understood to be speaking of biscuits on the sideboard

on such an understanding of there being some. (Again there is the question how circumstances may make this so.) So we may see Austin's deepest and most central concern as *truth*, the rest just flowing from what he found there.

This is pretty much the core idea of occasion-sensitivity. One specifies some things to speak of—being blue, say, the sky—such that in speaking of them in a certain structured way—saying the sky to be blue—one *might*, if things go well, say *something* to be so. Then in speaking of those things, in that way, and saying *something* to be so, one might say any of many distinguishable things. I mean the statement to be as general as just stated. Or rather, I mean it to hold for any *sublunary* things to speak of. For the nonce, I leave mathematics to one side. The core idea, naturally enough, has a corollary for semantics—by which I mean here a theory of what expressions of a language mean. Expressions of a language identify things to talk about—as 'being blue' identifies being blue, 'The North Sea' identifies the North Sea, and 'The North Sea is blue' identifies the North Sea's being blue. By the thesis, in talking about those things (in a given structured way) one might say any of many things. So what the expressions *mean* cannot fix some one of these as that which is thus said. So what they mean cannot fix any one condition as the condition for 'their' truth. Meaning cannot connect to truth like that.

I have encountered sheer incomprehension of the thought that meaning and truth do not connect like that, as if the very notion *meaning* (or what we choose to call that) excludes things being otherwise. I think that reaction is symptomatic: in the first place, of an inability to think of an expression's meaning what it does as an autonomous circumstance within the fabric of the various ways the world is, an a priori conviction that an expression's meaning something, if this is anything, must be something else. Which, in turn, is symptomatic of residual empiricism: the idea of privileged facts, whose status as such is unproblematic; and the rest of reality which, if really that at all, must be constructible out of the thus privileged. And (for reasons I have tried to bring out elsewhere) a conviction that we know before we look which circumstances have which status. Again, the only antidote to that frame of mind I know of is to take Chomsky seriously: before one asks what *else* some phenomenon might be, first ask what it actually looks like—how its various instances are liable to differ from one another, so, at least in that respect, how they are apt for being marked as the ones they are. Often a careful look in this direction will simply remove the urge to look for something else for the phenomenon to be. As Austin's careful look at the present one should do in the case of expressions meaning what they do.

If I speak of the sky as blue, I may, in doing that, speak of any of many things. So, if, on an occasion, I say, 'The sky is blue', thus speaking of the sky as blue, and *if* I thus say something in particular, the circumstances of my speaking must determine what this is. Meanings alone cannot do this. That was Austin's just-cited idea. How many ways can circumstance matter to what is thus said?

One can see this as a central question in a justly famous symposium in 1950, featuring J. L. Austin and P. F. Strawson on the topic *truth*. In it, Austin said,

A statement is made and its making is an historic event, the utterance by a certain speaker or writer of certain words (a sentence) to an audience with reference to an historic situation, event or what not. (Austin 1979/1950, pp. 119–20)

Strawson took exception. (See Strawson 1950.) Rightly sensitive to (one part of) grammar, he insisted that ‘The statement that P’ (scare quotes) is something there *is* to be made, detachable from any making of it by any particular speaker on any particular occasion. (This echoes Frege’s insistence that a judgement—what is judged—is always detachable from any judging of it. (Cf., e.g., Frege 1979/1915, p. 251.))

But the point for Austin is not whether you *can* talk (correctly) in the way Strawson suggests. And it is not, or should not be, for Strawson whether you can, equally, talk in the way Austin suggests. There is a question here of what idea of a truth-bearer to take as fundamental in an attempt to understand what truth is. And a correlative question as to just what one would be saying in speaking in the form Strawson recommends.

That issue might be put as follows. We might conceive of a syntactic theory of English as consisting of a finite vocabulary, a finite set of rules for constructing structures out of this vocabulary (or structures of it), and some definition of *whole structure*. The theory would then generate, by application of its rules on its vocabulary, a certain set of whole structures. A very minimal requirement on the correctness of a theory so conceived would then be: for any English sentence there should be some unique, or uniquely designated, whole structure which describes it (on some intended notion of describing); for any two different sentences there should be two different such structures; for any two different such structures, there should be two different sentences they describe. (I assume here that different structures belong to different sentences.)

Parallel to this one might think of a theory which generated all the different statements there were to be made (precisely) in speaking of the sky, and speaking of it as blue, or, say, of something in speaking of it as blue, and similarly for other sublunary topics one might speak of. The theory’s minimal goal would be to generate specifications of statements such that for each there is exactly one statement it fit; and for each statement, exactly one specification which fit it. That a given specification fit a given statement (say, the one someone made on some occasion) would then answer, precisely and unequivocally, the question *which* statement that was, construing ‘statement’ in Strawson’s way.

Does this goal make sense? If, but only if, there is some right way of *counting* ways of understanding given words, or what they speak of—understandings, say, of something being blue—when these understandings are detached from some particular words that bear them, are considered as understandings *for* words to bear—as, for Frege, the content of a judgement, a thought, is always detachable

from any judging of it. That is, if, but only if, there is some specifiable set of features of an understanding which are just those which *can* make a difference to whether we have to do with some given understanding, or some different one. In that case—but only then—the work of circumstances would be essentially that of disambiguation—choosing one item from some well-determined range of alternatives. And in that case, Austin's choice of fundamental truth-bearer would be otiose.

Conversely, Austin made the choice he did in the conviction that the idea of a generative theory of understandings *for* words to bear was just fantasy. Here is a way to think of that. For any historical stating of something in given words, there is a potentially indefinite variety of independent questions that *might* be posed, or arise, as to how that stating, or what it spoke of (a cat having mange, say), was to be understood. One might look at particular cats in a host of particular conditions, and ask of each whether being as that cat is, is having mange on the understanding on which that stating spoke of this. If circumstances of a stating matter to what was stated along the lines indicated here—lines on which what *was* stated is fixed by what one then had a right to expect of things being as stated (of a cat's having mange, say)—then there is in principle no end of opportunities for circumstances of a stating to matter to what was stated. There is no point at which circumstances choose for us some truth-evaluable item which is itself immune in principle to admitting of different further understandings—no point at which, through appeal to circumstance, we arrive at the sort of invisible, intangible truth-bearer (what Frege called a 'thought') which, Frege held, was the *only* thing that could really make a determinate question of truth arise. This is the outcome for which Austin's choice of truth-bearer was designed to make provision. If it is how things are, then, reading Frege as above, he had the wrong idea of what a determinate question of truth would be.

So on one reading of Frege (I think there is another, but one must work for it), Strawson was simply being a good Fregean. I elaborate. In 'Der Gedanke' (1918–19), just after its first two opening paragraphs, Frege starts off on a quest for that which 'actually raises a question of truth'—what he will regard as a truth-bearer. He begins with pictures, conceived as pure visible, tangible things—a picture as a canvas, covered with paint in a certain way. He points out that such a thing can only represent something as so—so only raise a question of truth—if an (adequate) intention attaches to it. A picture (so conceived) which might be a picture of Kölner Dom (perhaps there is some resemblance) might also be many other things—a fantasy, or a genre picture, or etc. And if it does depict Kölner Dom, still, it might be representing any of many things as so, depending on the manner of representation it adopts. A blue patch in an image of a wall, for example, might represent shadow at a certain time of day, or it might represent a blue patch on a certain wall. What might be taken in any of many ways, as representing any of many things as so, does not, for Frege, represent a definite question of truth. Since whether things are as it represents

depends on in which of these ways it represents things (if there is such a thing as ‘in which way it represents things’), it cannot yet be decided whether it is true or false. Contraposing, if something did raise a question of truth, then it would be something which could not represent in any of many ways, but which essentially represented in just one way, a way *not* admitting of understandings.

Such are Frege’s truth-bearers. They relate to circumstance in a certain way. If, in given words, I express such a truth-bearer, then, of course, it may be up to circumstance to fix what this truth-bearer was. But once it is fixed what it is, it cannot be up to circumstance to decide when *it* would be true. There is just nothing to decide. So the work of circumstance cancels out in this way at the point where we arrive at a truth-bearer. Suppose we make a further assumption: for any such truth-bearer, there are properties (ones, say, of representing in *this* way, of representing in *that*) which identify it as the truth-bearer that it is independent of any occasion of expressing, or producing, it; and wherever one expresses something with just those properties, one has expressed that truth-bearer. So one has expressed *a* truth-bearer. There is then no more work for circumstance to do in deciding when anything would be true. At this point we arrive at the picture which Austin rejects, and which Strawson—perhaps not quite realizing it—insists on. It is an interesting result, I think, that a truth-bearer, so conceived, is *not* what it looks like for a definite question of truth to have been raised.

When I said, say, ‘The room is dark’, my words were rightly understood in a certain way. Here, as elsewhere in philosophy, the locution ‘a certain way’, and relatives, may mislead in a certain way. For it suggests, where it should not, some determinate domain of discrete ways, to be counted according to some one effective principle for counting them. When Pia said ‘Sid grunts’, she spoke on a certain understanding of being a grunter. When Sid said ‘The room is dark’, he spoke on a certain understanding of a room being dark. ‘What understanding?’, one might be inclined to ask. Which may seem to call for an answer of the form, ‘The understanding such that . . .’, where what filled the blank would uniquely identify some *one* understanding there is to have of being a grunter, or of a room being dark. ‘It is that understanding on which someone would count as a grunter just in case (he were) such-and-such.’ *That* is what an answer would look like, on one understanding of ‘a certain understanding’, or of ‘understanding in a certain way’. Here, the question ‘What understanding?’ is read as a question ‘Which understanding?’ would be, with that ‘which’ referring to no contextually fixed special range of choices.

There is another way of understanding ‘a certain way’. On it, understanding being a grunter in a certain way would just be understanding it as it was to be understood on Pia’s speaking of it—understanding it so as to be that of which she spoke. When would one be doing that? Just when one was understanding it as it ought to have been understood in the circumstances of her speaking of it. There is determinacy in what one would have to do to be doing that. In the



circumstances, some things would be misunderstandings, others would not. It would be a misunderstanding, say, to think that anyone would be a grunter on the understanding on which Pia spoke of this if, hit hard enough in the solar plexus, he would grunt. One understands Pia's words in the right way—as they ought to be understood—if one does not lapse into misunderstanding. What one would have to do to do that, on the present story of a truth-bearer, is fixed by nothing less than the significance of the circumstances in which she spoke: that is, by their having the significance they did. Fine. But *which* understanding is that certain one on which Pia spoke of being a grunter? On the view I recommend, this last question is ill-formed. It has no answer. One cannot count understandings like that. Nor *need* the locution 'a certain understanding' suggest otherwise.

I offered this as an illustration. The locution crops up elsewhere in philosophy. 'You should have seen the way Sid's Lexus was when the police found it.' 'You should have seen the way Pia looked when Sid told her the news.' One *might* ask, 'Which way?', expecting, as an answer, 'Such-and-such way'. But no such answer need be in the offing. *Perhaps* whenever Sid looks at his Lexus, it may be said that it looks a certain way to him. (Though this sounds like a typically philosophical thing to say, and may be riding roughshod over occasion-sensitivities in notions like 'a way it looked' which are there for good reason.) But then this need come to no more than that some looking-to-him was going on. Exactly when would things, or the (or a) Lexus, look to Sid precisely as they, or it, then did? Well, Sid's Lexus, looking as it did, did that. There may be nothing more to say as to what *else* might.

The Lexus may have been such that such-and-such. That is a way for *a* thing, or at least a Lexus, to be. There is an indefinitely large range of distinct cases in which a Lexus, in being as it was, would be that way. Such a way for a thing to be belongs to the conceptual. A concept, if satisfied, *could* have been satisfied even if things, and what satisfies it, were not *just* the way they are. Even if the concept is of being Frege, so could be satisfied only by him, he could have satisfied it in being other than he is, and even had the *Lusitania* remained afloat. There is a *range* of cases of something being such as to satisfy it, and a determinate demand on falling within that range. So an answer of that form would identify something general: a condition of the Lexus which *could* be found in other circumstances, perhaps in other cars, if only such-and-such. Whereas if the Lexus' being a certain way just comes to its being as it is, then there is no such generality, no range of cases. To be in that condition, nothing will do but being *that* Lexus as it then was. There is the way we were back then. But that way need not be being thus and so. Similarly, where there is a way something looks, that way need not be looking thus and so.

To finish the picture of occasion-sensitivity as such, it is time to speak of nonsense. 'Down pub the he went', 'Milk me sugar', 'All mimsey were the borogroves', 'Going to Grantchester went to Grantchester', 'He is more identical

than I, 'The length of my bed has chocolate undertones'. These are all specimens of nonsense. Now consider the sorts of things that give philosophy a bad name. 'The sky is blue or it isn't.' 'What do you mean?' (With Wittgenstein's interlocutor) 'You know what those words mean, don't you? Well, I am using them in the sense you are familiar with.' Or, 'I *know* I had ten toes when I left the house this morning.' 'What do you mean?' 'Well, I know it or I don't, right? And I say I do.' The sky is blue on some ways of understanding its so being, and not on others. There is nothing *else* it would be for it to be blue or not. So if you manage, in calling the sky blue, so to speak that there is nothing but what being blue is, and what the sky is, as such to fix when things would be as you said, then you simply have not managed to say anything to be so, either truly or falsely. This, too, is a form of nonsense; one that some philosophers have positively aimed at.

What you say (if anything) in describing things in given terms always depends on the circumstances of your saying it. For you to have made good enough sense to have said something either true or false, circumstances must do work which they *can* always fail at. On a sunny day, someone, out of the blue, may call the sky blue. There is a truth near to mind that could be so expressed. We may count him as having expressed it. If, out of the blue, someone tells us that Sid is blue, we are likely to be baffled. When we encounter his blue-tinged, rather troglodytic complexion, we may see how one *could* call that someone being blue, and then may be willing to allow that *that* is things being as said. If, pointing to a thoroughly overcast sky, someone sighted says, out of the blue, 'The sky is blue', there may be *no* answer to the question what he said, or none which settled how things would be if he were right—even if there *are*, as there are, *some* truths that could sometimes be told in so describing an overcast sky.

So it goes with (nearly) ordinary talk. But it not infrequently happens that for a philosopher to say what he aims to say, the circumstances of his speaking would have to be making *no* substantial contribution to fixing what that was. Only the concepts he deployed could do that work. Philosophy *can* seem, by its nature, to impose such a requirement. If, say, I aim to tell you what it is to do something intentionally, I would miss my target if I merely told you what it was to do so on some special understanding of *intentionally*, or, even worse (it may seem) on some special understanding of *what it is*. Thus Grice's concern to show that language obliges philosophy's seeming demands, and, in fact, makes those seemingly off-target answers otiose: what was done intentionally *was* so done, no matter what the occasion for asking (even if it may sometimes seem odd to say so).

But if occasion-sensitivity raises such spectres, it also, in the same stroke, offers means to banish them. One can, for one thing, reframe the question, asking what sorts of different understandings *intentionally* may bear, and why: how one such may differ from others; how different ones relate. One *might* do that by asking, rather than 'When would something be done intentionally' (the

question that occasion-sensitivity makes unanswerable in principle), the perhaps answerable question, 'When would we speak truth in saying someone to have done something intentionally?'. In any event, such spectres had better evaporate if there is to be such a thing as philosophy at all: there *is* occasion-sensitivity. A philosopher committed to saying when things would be thus and so, on *no* contingent understanding of things so being is thereby committed to talking nonsense.

Inspired by Grice (see 1989/1967), many have learned contempt for questions of the form 'When would we say such-and-such?', or 'How would we describe such-and-such?'. The contemptuous idea would run like this. Consider, say, a range of people—candidate knowers that it is now mid-December. Or, again, some range of (physical) objects—candidates, say, for being red or not. Now let us partition the range in two ways. First, we partition it into those cases of which 'we would say' that they knew it was December, or that the objects were red, and those cases of which we would not say this. Then we partition it into those cases of which it is *true* that they know it is December, or that it is red, and those in which it is *not* true (or false) to say this. There is no reason, the thought is, why these two partitionings should coincide. There are all sorts of reasons why we might or might not say something. Truth is another matter. And if we are interested in the concept of knowledge, or of being red, then *truth* is our concern.

To make the situation seem yet graver, we might portray it this way. What we would say when is a matter of what we are inclined to do; thus, of psychology. But whether the tomato on the sideboard is red is not a psychological matter. Nor is whether I know it is December (a standing in the 'space of reasons', even if there is some psychological state I must be in to enjoy it). So to allow psychology to bear on truth as it (supposedly) would if 'what we would say' bore on this is, the thought is, just to start off on the wrong foot.

Occasion-sensitivity makes for two errors at the start in this way of putting things. First, we cannot partition cases of things being coloured as they are into those which *as such* are ones of something being coloured red and those which *as such* are not. There is, in general, no such thing as being red, independent of any particular understanding of so being. Even less could there be such a partitioning in the case of knowing such-and-such, or trying to do such-and-such. Second, and by the same token, there is no such thing as the cases which are those we would call ones of something being red—apart from some particular occasion for the calling. Again, all the less for knowing something, or trying to do something. There is not yet any sensible notion of discord or harmony between these two (non-existent) partitionings.

There are, on the other hand, other readings for 'When would we say' questions. It would not be *too* much of a stretch to read them as asking when it would be true to say . . . ; that is, as when, say, someone would count as knowing that P, or trying to do such-and-such. If you knew what trying was (as 'we'

presumably do), what would you so count? Only now, with occasion-sensitivity in view, the ‘When’ in these questions refers not just to the condition of the imagined subject—when the candidate knower, or tryer, is thus and so, but also to the occasion for so saying. And the question ‘What would you (*qua* understander) so count?’ becomes ‘What would you so count *when?*’.

The reason for preferring such questions to *prima facie*, more straightforward questions as to when such-and-such would be the case—where there *is* such reason—is not that somehow the formal mode is more, well, formal, or more precise, or anything of that sort. It is just that the straightforward question is liable to *admit* no answer. When would it be so that someone saw an orange? Well, what are we to understand by *seeing* an orange (or seeing an *orange*)? The answer is to be gleaned neither from what seeing is, nor from what oranges are, as such.

One can, of course, try, with Grice, to hold out for the view that, really, language makes ‘when’s, read as just indicated, otiose: there is no ‘when’ about it: to be such as to have tried is to be so, no matter what the occasion for asking after this. One trouble with this, as I try to argue in the essays in Part I, is that it collides with the linguistic facts. But another is that, on all the evidence so far, it shackles us with the same dreary, and, in fact, insoluble, philosophical problems. Thompson Clarke (1972) provided one brilliant case study of this. One can find an equally good example in the case that first inspired Grice. (See Grice 1989/1961.) There is, or there sometimes counts as, a special way of viewing things, available only to me now. There are, on some occasions for their posing, intelligible questions as to how things looked so viewed, with true answers to the effect that they looked to me thus and so. Suppose that the truths one may thus *sometimes* tell are *always* to be told: if it is ever true to say things to have looked thus and so to me, then it is always true to say so. Now I see that armadillo before the cactus. There are ever so many ways for a cactus to look: *this* shade of green, *that* one, and so on. For each, there is the question whether the cactus then so looked to me: either it did, or it did not, full stop. There is then the question how my awareness of it so looking to me could relate to my awareness (if any) of it looking as it did, or *being* as it was. Here occasion-sensitivity banishes the question, and with that (if philosophy has shown anything clearly) a route of inquiry along which only madness lies.

In philosophy, as Austin remarked, it is no mean accomplishment to have said something intelligible enough to be clearly false. If we are occasion-sensitive thinkers (and speakers), this needs to be taken into account if we are to get even that far. The essays in Part II of this collection illustrate what that might come to. I will conclude by explaining their motivations as I recall them. Those motivations still move me, even though, were I to begin anew, each of these essays would, no doubt, sound considerably different. First, ‘Are Belief Ascriptions Opaque?’ (Chapter 8). Here two factors were at work. First, it seemed (and seems) to me that the phenomenon of opacity is widely misunderstood.

Notably, opacity, or intensionality, does not signal the possibility of expressing truths of the non-existent, whatever it might be to do that. I thought, and still think, that such misunderstandings as to what opacity is are one of the bad effects of a penchant for casting philosophical problems as linguistic ones. There are cases where linguistic formulations can expose a phenomenon which is otherwise easily missed—occasion-sensitivity being a notable case in point. But there are others where focusing on the linguistic traces of a phenomenon can distort its underlying nature. Intensionality, and intentionality, it seems to me, are a case in point. Second, it seemed to me that missing out on the occasion-sensitivity of such things as belief ascriptions worked to encourage misimpressions as to what opacity in fact comes to. This essay was an attempt to bring out just how that may be.

‘Vagueness, Observation, and Sorites’ (Chapter 9). It had been suggested, notably by Michael Dummett, that paradoxes such as the Sorites might signal that natural languages were inconsistent. I admit that I do not understand what it might be for a natural language to be inconsistent. Natural languages do not assert anything; *a fortiori*, do not assert inconsistent things. There is, of course, a long-standing idea, with many incarnations, according to which a natural language, in arranging for its expressions to mean what they do, commits to the truth of certain propositions, and may commit inconsistently, or at least in ways inconsistent with the facts. For example, perhaps the expression ‘straight line’ commits English to the existence of paths such that if A and B are any points on them, then, for any point C, the length from A to B is less than or equal to the distance from A to C plus that from C to B, and, second, such that if two of them are both orthogonal to some given line, then they never meet. But (perhaps) there are no such paths in space. So a crude version of the idea goes. But, to say the least, it is anything but clear that for expressions to mean what they do is for them to make such bald commitments, or that an expression, in meaning what it does, *could* make flatly inconsistent ones—that the notion of an expression meaning something works like that. It is, anyway, quite a long leap from the existence of paradoxes like the Sorites, or the Liar, to such a conclusion, and a leap which seemed, and seems, to me utterly unjustified. Again, it seemed to me that ignoring occasion-sensitivity has encouraged taking such a leap. The purpose of this essay was to try to say how.

‘Attitudes As States’ (Chapter 10). This was a first attempt at ideas which, developed further, were published—quite a few years later—as my book, *Unshadowed Thought* (2000). I tried to identify deep roots, in Fregean thought, for an idea that may seem to be just an artefact of trying to force philosophy into the mould of cognitive science (or, conversely, to represent cognitive science as philosophy). The idea I had here was that the deep root of a widespread, but hopeless, view of propositional attitudes was an assumption, which I then saw in Frege, to the effect that one could treat ‘thought’, or whatever word one chose to designate (grammatical) objects of such attitudes,

perfectly straightforwardly, and with no qualms or reservations, as a count noun. Whereas with occasion-sensitivity in the picture, we can allow the idea that thoughts are countable *some* play without taking it perfectly seriously, or at least without supposing that there is some particular way which is *the* way in which thoughts are to be counted. I should say that over the years I have become considerably more sympathetic to Frege than I was in 1989. So I would now be more inclined to put the matter this way: there is a nefarious idea as to the countability of objects of attitudes which many people have, no doubt, felt inspired by Frege to hold. How attached Frege himself was to that idea is open to debate. Anyway, if he did hold it, as he may have, it is an optional add-on to an underlying view of propositional attitudes (in particular, of judgments) which was exactly right.

‘On Concepts of Objects’ (Chapter 11). This is perhaps my favourite of the essays in this collection, though at the time it found no sympathetic audience. I think that this is perhaps because it is only a subtle step away from an absolutely crazy view. The inspiration for this essay is, of course, David Wiggins, who, quite some time ago, exposed the idea that the identity (or not) of an object, A, and an object, B, might be *relative* to something—e.g., to some sort of thing A was and B was. I do not dispute Wiggins’s result on that point. It is a *result*. But relativity is one thing, occasion-sensitivity another. One might, I think, make the right point here by speaking in terms of individual concepts. If we think of a concept as, essentially, of being such-and-such—so that wherever there is such a thing as being F, there is, accordingly, a concept which is precisely of being *that*—then we will find, fitting this formula, such things as the concept of being Frege, or the concept of being New College. Such things are perfectly good concepts, though with one distinctive feature. The concept of being Frege, as opposed to the concept of being bearded, is, of course, such that no one *other than* Frege could fit it, or have fitted it, no matter how things were (and accordingly such that there would have been no such concept but for the happy accident of Frege’s birth). All of this leaves it quite open whether, when one arrives at Frege’s deathbed an hour after the fact, what one encounters is Frege. And it may be that nothing in what it is for someone to be Frege delivers, as such, any unique answer to that question. In which case there will be various ways of thinking, and speaking, of Frege, and, corresponding to each, different sorts of truths to express as to what (or who) is, and what is not, Frege. This essay develops that thought.

‘On Constraints of Generality’ (Chapter 12). This essay was, of course, inspired by Gareth Evans. Again, Evans has a core idea that I would not wish to dispute (at least in its main and most important applications). To know what pain is (to have the concept pain) is to grasp what it would be for *one* to be in pain. And similarly for any psychological phenomenon. That good point is at the heart of Evans’s story about what it is to grasp a concept. But I thought, and think, that Evans over-generalizes the point. On his story (modulo views about possible

categorial structures in systems of concepts), to grasp any concept is to grasp, for any object one can think of, what it would be for that object to fall under that concept. I found two reasons for wanting to resist this story.

One is a basic point of Wittgenstein's methodology: we understand how concepts would apply in, or to, particular circumstances for which they are designed; a main source of philosophical (pseudo) problems is the impression that we understand their application—that it is 'just the same as what we are familiar with'—to circumstances for which they are really not defined at all. Wittgenstein's model of this is: I understand what it would be for it to be 5 p.m. here in London. I understand what it would be for it to be five hours earlier in Boston. But, for all that, I simply do not know what it would mean if someone said that when it was 5 p.m. here it was such-and-such time on the sun. Not that the sun is, intrinsically, the wrong sort of object to have times. Just that it is undefined what it would be for it to be such-and-such time on the sun. So, again, I understand what it would be for those green leaves on the tree to look darker in the early morning light than they do at noon. I may come to understand what it would be for green to look darker to Sid, in his condition, than it normally does. (Though this is something I would have to come to understand by seeing what particular condition of Sid is meant to make that so.) But I do not in general understand what it would be for green to look different to Sid than it does to Pia. It does not follow that this has a clear sense merely because those other things, just mentioned, do.

The second reason, which I go some way towards explaining in the essay, is that the idea of the generality constraint, as Evans develops it, seems rooted in a Tractarian idea of what it is for a thought, or proposition, to have elements, so to be structured, and why any proposition must be structured, with a particular sort of structure, in some one particular way. Later Wittgenstein repudiated that idea, and gave very good and clear reasons for doing so. I am even more convinced today than I was in the mid-1990s that this aspect of the case has led to ill effects on subsequent philosophy. In any event, this essay is an attempt to develop the inconsistency of Evans's idea with the occasion-sensitivity of the concepts in terms of which we in fact think and speak.

'A Sense of Occasion' (Chapter 13). Much that currently goes under the name of contextualism about knowledge I want no truck with. For it misses the basic motivation for contextualism in that case. Nonetheless, it seems to me (as it also seemed to Thompson Clarke) that there is no possibility of making acceptable sense out of the notion of knowledge without occasion-sensitivity very conspicuously in the picture. The inspiration for this essay was, in fact, John McDowell, who I see as at our end of a long, continuous, Oxford tradition. I believe he does not see himself in this way. This essay was an attempt to persuade him.

The tradition begins with John Cook Wilson at the beginning of the twentieth century. In brief, Cook Wilson's model for all genuine knowledge was knowledge in (very simple) mathematics. Think of a proof of, say, the proposition that there is no largest prime. With that proof firmly in mind, think of how you stand towards the proposition that there is no largest prime. If the proof is, in fact, proof for you, then you see how it is that there is no largest prime. What you see excludes all possibility of it not being so that there is none. Knowing there is no largest prime is, on one conception of knowledge, appreciating that (and how) there is no such possibility.

The main problem in epistemology, at least from the seventeenth century on, has always been how, on such a conception of knowledge, one could know such things as that one's steak was charred, or that the pig was in the pen. Descartes seized on one strategy for defusing the problem: one could postulate senses of *know*. This would be to give up on the idea of knowing that one's keys were in one's pocket on the same notion of knowing on which one could know that there was no largest prime, but to make that sound anodyne with the assurance that, in some weaker sense of 'know', one *could* know such things. One could not, perhaps, have *strict* knowledge as to the whereabouts of one's keys. But one could have what Descartes called 'moral certainty'. And that would do for our mundane human purposes. Locke (1959/1690, p. 332) expresses that idea in these words: 'The certainty of things existing *in rerum natura*, when we have the testimony of our senses for it, is not only as great as our frame can attain to, but as our condition needs'—an idea which Austin exposes, in *Sense and Sensibilia* (1962*b*), as a bad one, and which McDowell has exposed, conclusively to my mind, as a strategy that simply will not do. Many current versions of contextualism about knowledge simply elaborate it.

McDowell rightly insists that seeing the peccary before me on the trail *can* be, for me, proof that there is a peccary on the trail on just the same notion of proof on which a proof that there is no largest prime, when I grasp it, may be proof for me that there is no largest prime. Here he departs from Cook Wilson, and his disciple H. A. Prichard. The question is what permits this departure. The answer I suggest (as Austin did before me) is: occasion-sensitivity. But only when that idea is applied in the right way. Special features of knowledge call for some subtlety in its application to that case. This last essay in the present collection is an attempt to work them out.

A final word. Formulating issues in linguistic terms can clarify, but also, sometimes, obscure. If we want to understand what some feature of mental life is, we must ask after what we *say* of someone in attributing it to his. Occasion-sensitivity, among other things, makes this mandatory. But recognizing it is for getting *phenomena* in view. There is no *general* reason to model *them*, or their workings, on the language in which we represent them, or the workings by



which it achieves the representing it does. As, over several decades, I worked through the issues treated in these essays, I became more and more convinced of this—which, over time, drew me closer to Frege. Reading him as I now do, I would be much more cautious about attributing to him some of the views I have below. On the merits of the views attributed, my views have not changed.

All the work in this collection has benefited systematically from the profound influence on me of a few friends and (often) critics. I am especially indebted to John Campbell, Peter Sullivan, John McDowell, and Mike Martin.

# Part I

## Occasion-Sensitivity

*This page intentionally left blank*

# 1

## On What Is Strictly Speaking True

Saying, it goes without saying, is but one way of putting things into words. Among others are hinting at, implying, suggesting, and letting slip out, to name a distinguished few. But *how* is saying distinguished from such rival modes of conveying, or at least expressing things, and why does this matter? One answer to these twin questions was proposed some years ago by H. P. Grice,<sup>1</sup> and has enjoyed, ever since, a considerable vogue. What is said, Grice thinks, is often simpler than it seems. And he offers, in effect, a way to argue that it is wherever complexities may seem to be rearing their heads. Given the seriousness with which this offer has been taken up, and the relentlessness with which it has sometimes been applied, and given the resultant dismissal of phenomena better worried about, it is perhaps worth showing why Grice's suggestion cannot be correct. In making it, Grice offers up a bounty of things to dispute, but, for the sake of order, I will confine myself here to a principal few. These I will first sketch, under six distinct headings.

### A. The Saying–Implicating Dichotomy

The distinction between saying and otherwise putting things into words—whatever this may come to—is supposed to correspond neatly to a distinction between *what* is said and *what* is otherwise to be understood or gathered from what is said, or it's saying: If, strictly speaking, it wasn't said *that* P, in speaking words *w*, then P isn't part of what was said. P's being so is not, e.g., part of what was said to be so in the words *w*. Most important, it is not part of what must be so where what was said to be so in *w* is so. Thus, the test whether it can be said, correctly, to have been said that P does the extra duty, according to Grice, of showing whether what P describes as holding is required for the being so of what was said to be.

If it wasn't said *that* P in *w*, but P nevertheless would be to be understood in or by *w*, then P must have been put otherwise into those words, and thus falls on the other side of any 'said'–'something else' dichotomy. For the something

<sup>1</sup> Specifically, in his William James lectures, delivered at Harvard in 1967.

else in this case, Grice has a technical term: ‘implicated’.<sup>2</sup> Given the above view, what *is* implicated, and what implication is, may be specified by contrasting it with what is said, given a grasp of what what is said is. On the basis of a single example of drawing the distinction, Grice does just this:

I think it is clear that whatever B implied, suggested, meant, etc., in this example, is distinct from what B said which was simply that C had not been to prison yet, I wish to introduce, as a term of art, the verb *implicate* and the related nouns *implicature* (cf. *implying*) and *implication* (cf. what is implied). The point of this maneuver is to avoid having, on each occasion, to choose between this or that member of the family of verbs for which *implicate* is to do general duty. I shall, for the time being, at least, have to assume to a considerable extent an intuitive understanding of the meaning of *say* in such contexts, and an ability to recognize particular verbs as members of the family with which *implicate* is associated. (Grice 1975, pp. 43–4)

## B. The Autonomy of What is Said

‘In some cases’, Grice remarks, ‘the conventional meanings of the words used will determine what is implicated, besides helping to determine what is said’ (ibid.). But there are other cases as well, which Grice terms cases of ‘conversational implicature’. Assuming a notion of implicature in general, about which he has said little more than the above, Grice defines this particular kind of implicature as follows:

A man who, by (in, when) saying or (making as if to say) that p has implicated that q, may be said to have conversationally implicated that q, *provided that* (1) he is to be presumed to be observing the conversational maxims, or at least the cooperative principle;<sup>3</sup> (2) the supposition that he is aware that, or thinks that, q is required in order to make his saying or making as if to say p (or doing so in *those* terms) consistent with this presumption; and (3) the speaker thinks<sup>4</sup> (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (2) is required. (p. 44)

<sup>2</sup> This class of ways of not saying things is apparently meant to exhaust the ways of not saying things by (or while) putting them in one’s words, i.e., while they contribute to ‘the total signification of an utterance’ (as I read it, how the words as spoken, or perhaps the speaking of those words, are/is to be understood). For, in his third William James lecture, Grice says, ‘The total signification of an utterance may be regarded as divisible in two different ways: first, one may distinguish within the total signification, between what is said (in a favoured sense) and what is implicated; and second, one may distinguish between what *is* part of the conventional force (or meaning) of the utterance and what is not. This yields three possible elements: what is said, what is conventionally implicated, and what is non-conventionally implicated.’

<sup>3</sup> At this point in *our* story, it doesn’t matter what these are. Just suppose them to be some clear and definite principles or other.

<sup>4</sup> This sort of reference to the speaker’s internal mental states is enough by itself, I think, to prevent conversational implicatures from having anything like the sort of theoretical interest which Grice thinks they do. But *that* is one of the quirks in Grice that I will *not* take up here.

One key fact about conversational implicature is that, on Grice's view, nothing *can* be conversationally implicated unless something *distinct* from it is said.

Again in Grice's words,

The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a *conversational* implicature; it will be a conventional implicature . . . .

. . . A general pattern for the working out of a conversational implicature might be given as follows: 'He has said that p; there is no reason to suppose that he is not observing the maxims, or at least the CP; he could not be doing this unless he thought that q; . . . (further steps) . . . and so he has implicated that q. (ibid.)

Whatever it was that was said in words w, then, one must be able to understand (and hence to say) *what it is that was said without understanding (or saying) what it was that was implicated. If P was said, and Q was implicated, then that Q would be, or was to be understood cannot be part of what one needs to know in order to know that it was P that was said.* Abstracting from psychological attitudes, the point may be put like this: If P and P' are two distinct things there are to be said in some words or other, and if Q is something that might be implicated, then whether it was P or P' that was said in words w cannot depend on whether Q was implicated or not. And so, the features of what was to be understood in or by w which would be present if Q was implicated, but not if Q was not, cannot be features needed to individuate P—e.g., to distinguish it from P'. For, unless it could be known that it was P that was said without recognizing the presence of such features, or understanding that Q was implicated, one could not *work out* on the basis of P's having been said that Q was implicated.

### C. On Saying

Though Grice assumes considerable intuitive understanding of what saying is, at least two things are clear about what *he* means (or thinks) saying of the relevant sort to be. First, by his lights, saying is something such that *what* is said is the sort of thing which *may* be either true or false—among many other evaluations to which it may be subject (e.g., plausible/implausible, absurd/sensible, too obvious for words/interesting, what the speaker couldn't seriously believe/what he could)—all of which are liable to enter into the working out of what is implicated on the basis of what is said, and do so quite regularly in Grice's examples.

Second, however, Grice also intends what is said 'to be closely related to the conventional meaning of the words (the sentence) . . . uttered' (p. 44)—so closely, in fact, that it is fixed fully by these meanings, 'together with the identity of any references that may be involved' (p. 50) (where by 'reference' Grice clearly means *referent!*).

Thus, what is said in a speaking of given words, in speaking a given language, is determined by what those words mean in that language (on that speaking, if they are ambiguous in the language) plus other factors of a narrowly specified sort. What this means will be set out more exactly in due time. Grice's own and lone example, however, should fix things clearly enough:

Suppose someone to have said, 'He's in the grip of a vice.' Given a knowledge of the English language, but no knowledge of the circumstances of utterance, one would know that he had said about some particular male . . . x, that at the time of utterance (whatever that was), either (1) x was unable to rid himself of a certain bad character trait or (2) some part of x's person was caught in a certain kind of tool . . . (approximate account, of course).

But for a full identification of what the speaker had 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 the utterance of the phrase *in the grip of a vice* (a decision between (1) and (2)). (p. 44)

#### D. How what is Said Makes its Presence Known

Largely, we detect that which was said in a given speaking of words simply through relying on our intuition about what was said (or, more pertinently, what *would* be said in speaking those English words).

We can simply ask ourselves, that is, whether we would say, for a given A, that a speaker or his words *said* that A. (Grice: 'If I say (smugly), *He is an Englishman; He is, therefore, brave.* . . . I do not want to say that I have *said* (in the favored sense) that it follows from his being an Englishman that he is brave . . . ' (p. 45).) But intuition of this sort may be joggled, at least, by intuition of another. Suppose, without yet positively identifying that which was said in words w, we identify something which would make those words strictly speaking true. Then suppose we find something else which was also conveyed by or in the words, but which might be false while (or even though) our something obtained. Then Grice takes that as good grounds for denying that this last item was part of what was *said*. For example, in the above case, Grice takes it that the relevant person's being an Englishman *and* being brave is enough to make what was said, strictly speaking, true—or at least not, strictly speaking, false. Since all that may very well be so, while it doesn't follow from his being an Englishman that he is brave, *that* the one thing follows from the other is not, according to Grice, part of what is said in the above words.

#### E. How What is Implicated Makes its Presence Known

Again, largely by intuition, and by contrast to what is said. Once we have fixed what was *said* in words w, anything else which would be (to be) understood

by those words must be implicated,<sup>5</sup> and *not* part of what was said. But here, too, there are tests by which to jog one's intuitions, at least for conversational implicatures. Since what is said is fixed (modulo referents) by the meanings of the words used, this will be a relatively stable feature across any and all speaking of those words. (This means, as we shall see, that there will always be *some* way of reporting what was said in words *w* which is valid and correct for every speaking of *w* (as meaning what they *do* mean in the relevant language).) By contrast, what is conversationally implicated in a given speaking of *w* is fixed by some features or other of that speaking—on Grice's account, properties of the speaker while producing it. Therefore, first, we may expect such an implicature to be a relatively unstable feature of speakings of *w*. If *q* is conversationally implicated in one speaking of *w*, it should be possible to find or devise *other* speakings where it is not. Grice terms this *cancellability*. Second, as long as the relevant features of a speaking remain fixed, it should be relatively difficult (though not universally impossible) to get rid of a given implicature merely by changing forms of words, while, in the relevant sense, continuing to *say*, strictly speaking, the same thing. This Grice terms 'non-detachability'. Finally, attached to the doctrines being considered here is a rough sort of 'theory' of how, due to what Grice terms 'conversational maxims', conversational implicatures could (or perhaps would) arise in the speaking of words. (The sense in which this might be a theory is the topic of the last section of this paper.)

If, for some part of the (correct) way of taking someone's words,<sup>6</sup> the theory can show how this could have come to be the way one *would* take them, then this is supposed by Grice to be at least strong evidence for the thesis that that part is conversationally implicated, and hence not strictly speaking said. Where such a 'demonstration' is available, and where the other features of a conversational implicature sketched above are present, this provides, according to Grice, at least 'a more or less strong *prima facie* case in favor of the presence of a conversational implicature'.<sup>7</sup> When it comes down to cases, Grice, to put it mildly, relies heavily on this assumption.

## F. Applications

Implicature is meant by Grice to undermine what he terms 'A-theses'. What *is* an A-thesis? There is, at any rate, a feature of certain philosophical theses which appears to catch Grice's eye, and to lead him to group them together under this heading. This feature is, roughly, the making of a claim concerning the limited

<sup>5</sup> At least if one of the verbs in the relevant class fits. I *think* this will always be the case, but since it is not *perfectly* clear what the class is, I add this escape clause.

<sup>6</sup> In Grice's terminology, 'part of the total signification of the utterance'.

<sup>7</sup> William James lectures, lecture 3.



applicability of some notion—for example, the notion of an action’s being done voluntarily, or of someone’s knowing that such and such, or of someone’s trying to do something or other (which he did do). The claim of an A-thesis, for some such notion, is that it may not be correct, or even sensible to apply that notion to an object of the relevant kind (a doing, say, or a knower and an appropriate object of knowledge), either by saying that the object *is* what the notion is a notion of being, or that it is not. Further, this will not be correct (or sensible) unless certain conditions are met.<sup>8</sup> So, for example, not every action will always be describable (correctly) either as having been done voluntarily, or as having not been so done, or as something the doer *tried* or didn’t try to do. And it will not be the case for just any knower and any object of knowledge that, on any occasion of describing, at least, the knower is correctly described either as knowing that object, or as not knowing it. So, for example, not every speaking of words which expresses such a notion and refers to such an object—not every speaking of the words, ‘It was/wasn’t done voluntarily’, for example, or, ‘x knows/doesn’t know that P’—will be such that in it something sensible was said, and in particular such that there is something which, in it, was said to be so, and which might either *be* so or not.

On some A-theses, some objects may, as such, be excluded from being correctly describable in either the one way or the other—e.g., as done voluntarily or not. A more interesting aspect of many A-theses, however, is that, on them, objects which may *sometimes* be describable in one way or another—e.g., as done voluntarily—will, on other occasions, not be so describable. What this will depend on is what describing of the one and the same object (doing, knowing, etc.) is in question. So the conditions to be met for the applicability of a notion are, in general, not just conditions on the objects in question, but also conditions on describings, or perhaps speakings in which the describing is done.

Now, what is the Gricean strategy for undermining such theses? Generally, in establishing this part of an A-thesis, one needs to show that when judging the truth or falsity of descriptions of the relevant kind—e.g., of an action as done voluntarily—what one judges to be so or not is such that nothing of that sort will have been said to be so unless certain conditions are met—in the same way that nothing will have been said, or said to be so in uttering the words ‘Bill is a friend’,

<sup>8</sup> The key point, on *my* reading of A-theses, is that, unless the extra conditions are fulfilled, one will not have said anything sensible at all—in Fregean terminology, one will not have succeeded in expressing a complete thought. As this suggests (though perhaps not *all* A-theorists would go along with the suggestion), one will not have said anything true, or anything false. Here, then, I will take it that ‘not correct to say’ implies, minimally, ‘not true to say’.

To emphasize the obvious: the point is *not* that there *is* something to be said, viz., ‘that X knows that p’ (‘did A voluntarily’, etc.), which, however, for one reason or another, it is not *correct* to say on some occasions. The point is not to be put, that is, by saying that it is not correct to say *that* X knows that p (or whatever). Rather, what it is not correct, under some circumstances, to say is, e.g., ‘X knows that P’, as if one *had* thereby succeeded in saying something that, in any case, there was to be said.

unless the conditions of that utterance are such that someone can correctly be said to *be* the Bill (or whoever) referred to in that occurrence of 'Bill'. And that requires that what one thus judges about *be* part of what is said, and said to be so in giving descriptions of the relevant kind. Now, suppose it could be shown that what is allegedly part of what is said *would* normally be part of what would be implicated, even were it *not* part of what was said. Then, *mere* evidence that what the A-thesis posits as part of what is said is normally part of what is to be *understood* in relevant speakings or describings would no longer be evidence in favour of the A-thesis, as opposed to some simpler rival on which, e.g., something either true or false will always have been said in describing *any* action as 'done voluntarily'. Since Grice takes it that this sort of evidence is all there is in favour of A-theses, he hopes, *via* the notion of implicature to remove any cogent reason for believing in them. What I hope to show a bit later is that mere evidence of the above sort, while it exists in abundance, is, *pace* Grice, not the only, and certainly not the most important reason one might have for believing in an A-thesis, and, in fact, that the real issues are quite poorly cast in these terms.

None of A–E are true, it seems to me, and the rest of this paper is an attempt to say why not. The central thesis which I will argue is that Grice's notion of saying is incoherent: There is *no* notion of saying satisfying the dual requirements on it that what is said may be true or false (and etc.), *and* that what is said is tied to the 'conventional meanings' of the words used in the way Grice envisages. So far, then, we have no idea where the boundary between saying and implicating goes. This is not necessarily disastrous for a notion of implicature. Surely *some* coherent notions of saying and what is said must be available. Once they are in, it is always open to define implicature contrastively in terms of them. Once we know what what is said is, we can always *define* implicature as, e.g., the rest of what is to be understood by or about a speaking of given words, whatever the 'rest' may turn out to be. But the way in which Grice fails to provide us either with what is said or with 'the rest', and specifically *what* he fails to envisage about what is said removes any force from the considerations he offers in particular cases in favour of something being implicated and not said: Or so I will try to show. This, I think, is important, at least if one cares about the *philosophical* issues at stake where A-theses have been offered.

### On What is Said

The next step is to demonstrate my central thesis. I will be brief, and regrettably stingy with data, having already indulged in greater expansiveness elsewhere (Travis 1975, chs. 1, 3; 1978; 1981*c*). Consider very simple sentences of subject–predicate form, for example, 'The kettle is black'. Suppose we want to know what is implicated in a speaking of this sentence. Then, of course, we must first know what was *said* in that speaking of those words. But what,

strictly speaking, *would* be said in such a speaking? Consider the following two contrasting cases:

- A. Max fills his shiny new aluminum kettle with the makings of a stew, and sets it over the campfire. An hour later, he informs Sam that he has done this. ‘That was pretty stupid’, Sam replies, and rushes out to the fire. He returns holding a soot-blackened pot and says (speaking truly), ‘Look. The kettle is black.’
- B. Everard and Clothilde are acquiring their first common batterie de cuisine. For many reasons, including tradition and presumed heat-retaining properties, they want only black pots. (Though what sort of black pot happens not to matter much.) Coincidentally, Max’s soot-blackened pot has come to rest precisely in the shopwindow into which they are now staring. Everard says, ‘Look. There’s a nice black pot.’ But Clothilde is more observant. ‘No it isn’t black’, she replies, ‘it’s only covered with soot. How careless of them to let *that* get in their window.’ And off they go elsewhere, with, to all appearances, Clothilde having spoken the truth.

It appears, then, that there are at least several *distinct* things to be said of a given kettle in one or another speaking of our sample words—among which, one thing which is *true* of a soot-blackened sample of aluminum, and another which is false of this same sample. If so, then it is incorrect to speak of ‘what would be said (strictly speaking) in speaking those words’—at least if this is to be understood as referring to one definite thing which there is to be said. Nor does the meaning of that English sentence enlighten us as to *which* thing was said in any given speaking, nor, consequently, how the implicatures of that speaking (if any) may be worked out.

Perhaps all this is *just* appearance. But whatever it is, it is an appearance which manifests itself quite regularly, for all sorts of predicates and along many dimensions. Sticking for a moment to the words ‘is black’, whether something counts as correctly described by them, on a speaking, may turn on an indefinite variety of issues besides whether soot coverings so count. Consider, for example, a ceramic-covered kettle with a cast iron core, versus a black ceramic-covered kettle with an aluminum core; a ‘black’ narcissus, with a green stem, a postage stamp which is black on one side, the sky at night (with its faint glow), an actor’s costume in a given stage light, something very shiny, black cod (white on the inside), or lumpfish caviar, which ‘bleeds’ a blue-green (to name a few potential problems). Each of these things may count as being black, depending on the way of counting. And corresponding to each of these ways is, presumably, a distinct thing to be said in *calling* something black (in the words ‘It’s black’, presumably, but if not in those, then in others).

Further, the appearances are not peculiar to ‘It’s black’ or to ‘colour predicates’. Consider, for example,

- a. 'This duck weighs 3 kg', said of: (i) a dressed duck which registers 3 kg on an accurate scale, but with 1 kg weight in its cavity; (ii) a live duck which will weigh in at 3 kg when dressed; (iii) a stuffed and mounted duck which registers 3 kg with the stuffings (with the mountings); (iv) a duck weighed during a space flight; (v) a duck weighed in a gravity simulator simulating 2 g; (vi) a duck packaged in plastic, weighing 3 kg with its neck, feet, and internal parts (with the plastic), but where these parts have been detached; (vii) a duck weighed while inflated with helium; (viii) a duck weighed just after being force fed. In all these cases, there is a true point and a false point to be made in saying the duck to weigh 3 kg, either of which points, it appears, might be made in speaking the words, 'this duck weighs 3 kg'.
- b. 'He's a sailor', said of: (i) a stock broker whose 10 m yacht and weekend jaunts set him off from his squash-playing colleagues; (ii) an ordinary seaman who, whenever at sea, shows himself to be miserably and perpetually 'at sea'; (iii) the executive chef, or the bandleader on a cruise ship; (iv) a retired sailor, or one who has changed professions; (v) someone with a 'natural aptitude', but who has never sailed.
- c. 'This is a wooden table', said of: (i) a table covered in wood-grained vinyl, (ii) but over a 'real' (scrap) wood interior; (iii) a table with wood planks for a top, but with metal legs (think, e.g., of situations where nonconductivity is prized); (iv) similarly, a table put together with bolts or screws; (v) a particle board table (vi) with particles made out of crushed walnut shell; (vii) a cross-section of tree trunk, sometimes used as a table; (viii) a plywood table; (ix) a table minus legs, hung on the wall and exhibited; (x) a table-like structure made out of wood shavings, and so on.
- d. 'This kettle is my grandmother's', as said of a kettle (i) given to the speaker years ago (ii) just before his grandmother's death; (iii) stolen by the speaker's grandmother just before leaving her native village sixty years ago; (iv) used consistently and exclusively by her, but technically only communal property; (v) said to the repossession agent who is at the door because grandmother is six weeks behind in her payments.

And, obviously, the same kinds of variations in what there is to be said may arise in general for artefact words (is it a chair?), natural kind words (is this (taken from a typical urban river) water?), measurement and shape words (you call *that* square? is it really 10 metres high?), and so on. If there are any *forms* of words which are proof against casting such appearances, it is, to understate things, difficult to see what they might be.

The question may arise, though, whether the sorts of speakings sketched and alluded to here *really* exhibit, for the most part, variation in what there is to be said, strictly speaking, in speaking a given form of words. For, one may have the suspicion in one or another case that what one would *call* true, or what one

would call false (e.g., either the true or the false point to be made about the kettle stolen long ago by the speaker's grandmother) isn't really what is said in given words, e.g., 'This kettle is my grandmother's', but rather, perhaps—to use Grice's term—something *implicated* in the saying of something else. For example, perhaps in a given case, the speaker makes it obvious to his audience that he could not *really* be saying or believe that the kettle 'belongs to his grandmother' (what is strictly speaking said), but rather—that being out of the question—that it bears some other relation to her, which somehow resembles belonging to her—such as being at her disposition, or being something no one is likely to take away from her, etc. Or, consider Sam's first remark about the soot-blackened kettle. Someone might feel, perhaps, that strictly speaking, the kettle is not black, so strictly speaking, what Sam *said* is not true, and the point which *is* true in virtue of the kettle's soot-blackened covering, not one which Sam made—at least not by saying it, though perhaps one which he implicated. And *that* he may have done by relying on the shared knowledge of speaker and hearer that, the kettle being made from shiny aluminum, it obviously could not be even mistaken for black. (Symptomatically, one could just as easily have the opposite intuition—that really the sooty kettle *is* black, and Sam spoke the truth in saying it to be.)

Perhaps, in fact, some such suspicions are correct in one case or another. Perhaps, from time to time, we *do* call something true (or false), seeming to be speaking of what was said in given words, but not really doing so. But much more than this must be so if it is not to be so that, strictly speaking, there are, for example, many distinct things actually to be said in calling a given kettle black (at given time), in saying of it, 'This kettle is black'—if, that is, the conclusions drawn above from initial impressions are not correct. For the claim was that cases A and B formed what might be called a *contrasting pair*. Since there is nothing special about the number two, let us speak of a contrasting n-tuple, and *mean* by that the following: a contrasting n-tuple for given words—e.g., the English, 'The kettle is black'—will be an (ordered) list of examples—speakings, or possible speakings of the words—in each of which, we would perhaps naively and overhastily, take a different thing to have been said (or at least a different point to have been made) from that said (or made in any of the others). We will, thus, say such things as, 'That's true', and 'That's false', and perhaps also give other assessments, such as, 'Correct', 'Right', 'You can't be serious', 'That's deceitful', 'It's just like/not a bit like what you said it to be', under different circumstances for each member of the list—at least when we are doing what we should be inclined to regard as correct.

For any contrasting n-tuple, e.g., for the words 'The kettle is black', what needs to be shown to save something like Grice's notion of saying is that, barring ambiguities which those words may bear in English,<sup>9</sup> and barring such things as

<sup>9</sup> Or, of course, in general, whatever language is being spoken in relevant speakings of the words in question.

shifts in what object(s), time(s), etc. are spoken about in the words, there is *one* thing which is strictly speaking said throughout all  $n$  cases (where that thing is something that may be true or may be false). Otherwise, what is said cannot be 'closely tied to conventional meanings' in the way in which Grice would have it.<sup>10</sup> But let  $n$  be 17, say. And fix references—e.g., let all 17 speakings concern one kettle at one time, in one condition. Then, naive or not, we have, at any rate, identified 17 distinct things to be said in some words or other (whether or not in those words which here led us to perceive the 17 things), and thus 17 distinct candidates for the post of 'what is (or what would be) strictly speaking said', in speaking those words—e.g., in saying, 'The kettle is black'. Grice's notion demands that at most one of these candidates be the right one—and, if none of them, then some other. And there must be some fact of the matter—something about what those words mean, or the uses to which, strictly speaking, they may be put, in speaking English, or etc.—in virtue of which it is so that this candidate represents what is *really* to be said in those words—a demonstrable or at least recognizable fact of the matter, if any such thesis is recognizably to obtain.

But establishing such a thesis is an obviously hopeless task. Consider Sam's kettle. Does what is *really* said in saying 'It's black' require that it be black, so to speak, through and through, as in case B, and then, if so, how permanently, and does it matter in what light, or under what conditions? Must it, e.g., be made of cast iron or the like? Or will a sooty surface do? Can a kettle really become black through soot? Would it matter if it glowed orange at  $1000^{\circ}\text{C}$ , or  $100^{\circ}$ ? or  $50^{\circ}$ ?  $25^{\circ}$ , but not  $24^{\circ}$  or  $26^{\circ}$ ? Clearly, there *are* no intuitions to throw into the breach here, because clearly nothing we know about English instructs us one way or the other on such issues. A reason for this is not hard to find. 'Black' is an English word for a certain colour, viz., black. That, plausibly, is what the meaning of that word tells us. And, if you know what colour that word means, then, plausibly, you *know* the meaning of that English word.<sup>11</sup> 'Is' functions in English, among other things, perhaps, as a device for forming predicates. And the predicate 'is black' functions, among other ways, for describing some thing(s) as being (coloured) black.

Things of the above sort are what knowledge of the meanings of English words typically tells us. The fact remains, however, that there are a number of ways of counting, or regarding things as being or not being black, and we freely avail ourselves of any and all of these, as called for. Sometimes a kettle, or a wall, a chair, a rug, or a golden retriever *does* count as being black if it is, e.g., painted black, and sometimes it does not. The meaning of the English words 'is black' does not tell us which of these ways is *really* the right way of counting things as black or not black, for the very good reason that none of these *is* 'the right way' nor is any other.

<sup>10</sup> And *needs* to have it for his arguments in particular cases to succeed.

<sup>11</sup> Barring, as usual, ambiguities.

There are, then, many distinct ways in which things may, on occasion, correctly be counted as being black or as not—many distinct things, that is, that being black may come to. Being black is what something sometimes properly counts as if composed of a black substance, except for its red enamel coating, and what other times a thing may count as given the condition of its surface, regardless of anything else. It is what the waters of the North Sea are sometimes correctly said to be (just look at them), and other times correctly said not to be (draw out a glass, or a bucketful, and look at *it*). But the *meaning* of the words ‘is black’ cannot choose between these various ways of taking things as black or not, if those words are to succeed in meaning *is black*. For, consider  $n$  such ways, and suppose it *did* select some of these, and not others, as ways of being as said to be in the words ‘is black’ (on that meaning). Each of these  $n$  ways is a way in which something may sometimes (properly) count as being black. Conversely, there is *one* thing which each of them is a way of counting *as*, and that is (what else?) *being black*. Being black, then, is something which there are at least  $n$  correct ways of counting things as being. Now consider being what the words ‘is black’ *mean* (or what, *in virtue of their meaning alone* they say something to be, or describe it as being). By hypothesis  $n$ - $m$ , of our  $n$  ways of taking something to be black or not are *not* ways of taking it to be as said to be in the words ‘is black’ (or to be that which those words mean). Being black is something which these  $n$ - $m$  ways are, on occasion, (correct) ways of determining something to be; being what the words ‘is black’ *mean* is not. Hence, being black is not what being what the words ‘is black’ mean is, or, being black is not what the words ‘is black’ mean (or mean being). Or, the words, ‘is black’, do not mean *is black*. Since the words ‘is black’ *do* mean *is black*, our hypothesis that their meaning selects between various correct ways of counting things as black (or taking them to be black) must be false.

What the words ‘is black’ mean *may* tell us what we describe something *as* in speaking them as bearing that meaning—viz., that those words describe as being black. But it is inconsistent with their meaning’s doing *that* job, that it also tell us what would or would not count as *being* as described in those words, much less what they are and aren’t truly spoken of. For it is inconsistent with their settling *such* issues that they mean *being black*. Now consider a particular speaking of, say, ‘It’s black’—e.g., in speaking of a particular kettle, *K*. And consider all the contrasting  $n$ -tuples with respect to those words, constructible out of occasions of, in one way or another, correctly taking something to be black, and (apparently at least) so reporting. These exhibit, in the way sketched above, a certain set of candidates for that which *was* said in this speaking. Unless *one* of these candidates *is*, in fact, what was said, nothing will have been said in that speaking which might be true, or be false, and hence nothing will have been conversationally implicated either—most notably, none of the things which *must* be implicated if the appearances which we have generated about what is said in such speakings are to be explained away systematically, as *mere* appearance. For,

suppose that what was said is *no* such candidate, but, e.g., something determined (modulo referents) by a proper understanding of the *meanings* of the words used alone. Then it cannot be something which would be shown to be true (or false) in virtue of the kettle's being in the sooty state it is actually in. For that would be something that the meaning of the words cannot show to have been said. It would, in fact, be one of those candidates which that meaning cannot choose between. But nor can what was said be something the truth or falsity of which is left open by the sooty condition of the kettle, but might be settled by something else. For again, this would be to select from among our candidates (e.g., so as to eliminate the previously mentioned one), again in a way that the meaning of those words cannot. But nothing true or false can have been said of the kettle unless it is something which *is* so or not in virtue of the condition the kettle is actually in. And something the truth of which is neither settled nor left open by the condition the kettle is in, is not something which *might* be true, or might be false at all.

Sometimes, perhaps, in saying 'It's black', we *do* fail to say anything either true or false. But such failure certainly cannot be the rule. And wherever we do not fail, what actually gets said must be one or another of the candidates which the above-mentioned contrasting n-tuples reveal. Which candidate it is in any given case must be fixed by a proper understanding of the words as spoken. And, given what meanings do, that is an understanding which must extend beyond anything which a proper understanding of the *meanings* of the words used might determine. Whatever further facts about some particular speaking make it the case that that which was said in it *is* what was said, independent of them no candidate has a privileged status *vis-à-vis* the other, with respect to what it is that *would* be said in those words (insofar as such a notion is defined). And that being so, we can conclude that, depending on what the further facts are, any such candidate might be that which *was* (or would be) said in one speaking of the words or another. There are, then, many distinct things to be *said* in speaking given words, in calling a given kettle at a given time black. And the fact that further facts of some sort about a speaking are required to fix how those words are to be understood cannot show, or even tend to show, that what is thereby fixed is *not* a part of what was *said* in those words.

What meanings of words don't do, but a proper understanding of a speaking of them does, contributes, both by what is omitted and by what is done, to words being usable as they are. On the one hand, it is a convenience to put it mildly, whenever, on a particular way of counting things, we take something to be black, to be able to *say* that that is what we take it to be—and to do so in saying 'It's black'. But that is something we could not do unless those words meant *is black*. That isn't what we would be *saying* in so speaking if to be what those words *mean* being were to be something essentially more or less than black. On the other hand, on a particular occasion where, in counting in a particular way, we take something to be black, there is something that we take it to be which is



different from what one might take it to be, in counting in a different way, in taking that same thing to be black. Where we take the kettle to be what counts as being black, in virtue of its sooty condition, we may not for a moment suppose that it is what one would take it to be, in taking it to be black, where something counts as being so in virtue of its composition, excluding enamel coatings and the like. If in speaking of something, in saying 'It's black', we are to be able to say to be so no more and no less than we *take* to be so, in or by so taking it, then we must, by so speaking, be able to say no more no less than *is* so, in virtue of the way things are, on the relevant way of counting or taking things. That is, we must be able to say the right selection out of the above-mentioned set of things there are to be said in calling something black. And, for us to be able to do this in general, *all* of these candidates must be things that are to be said in saying 'It's black'. We may, of course, for all that, sometimes *fail* to say what we think in so speaking, and manage to say something else instead. But fortunately, nothing about the meaning of those words, or about what there *is* to be said in them on that meaning, makes this either irremediably or unavoidably so.

There are several unfortunate misunderstandings here that I want to do my best to block. First, to reiterate, I am not claiming that, e.g., the *sentence* 'The kettle is black' is 17 ways ambiguous, or 39, or any ways ambiguous. What I *am* claiming is that there are many distinct *things to be said* in speaking those words, on one given meaning. The point is precisely that what there is to be said is *not* determined by the meanings of the words involved, nor even by their meanings on some reading, nor even by that plus the identities of referents. Second, I am *not* claiming that nothing about the 'circumstances of an utterance', or etc., determines in any given case such things as what being black is to be taken to be, or what it is to be taken to be to be as described in that speaking of 'is black'. If this *is* determined, which I suppose it is, then it is not unreasonable to assume that it is determined by *something*. Nor am I denying, at this point, at least, that such matters are determined by intentions of the speaker, or plausible reasonings about what he would be doing, or any other factor which Grice might wish to serve up. What I am insisting on is that *what* is determined, by whatever means, is *what it is that was said*—which item, that is, out of a large range of potential candidates for things to be said, and moreover, things to be literally and strictly speaking said in those words. And the reason for insisting on this, to reiterate further, is that without such things being determined, there simply *is* nothing to be either true or false; nor could there be any fixed condition for anything being either the one or the other.

So, for once, at least, things *do* turn out as they seem to be: Where, as with 'is black' (or, for that matter, any predicate) there are n-tuples of contrasting things we would perceive as having been said in various utterances of those words *on one meaning*, recognizing contrasting conditions for the truth of these various perceived objects, it is safe to conclude that those words, on the one meaning, may indeed make any of many distinct contributions to what is said in speaking

them. There are, in short, many distinct things to be said in saying 'is black', for example, and correspondingly many things that may thereby be said to be so (with interpretations of any other surrounding words fixed). Where what is said may be true or false, then, what was said in speaking given words cannot be revealed by their meanings (and referents) on that speaking alone.

Grice, of course, allows from the beginning that a 'full identification' of what was said in speaking of given words depends on what the words 'conventionally' mean, and on more facts about the speaking besides. From one point of view, it may seem that all we have shown so far is what some of this 'more besides' is. But if what has been shown is not (as I think it is) quite different from what Grice had in mind, it is at least quite different from what his discussions of saying and implicating require, as we shall soon see. Parenthetically, what *is* this 'more besides'? There are two views: on one view, the extra factor is how the words, as spoken (e.g., 'is black'), are to be understood—specifically, what is to be understood as being said in them (e.g., about something). On this view, the meanings of the words, together with a proper understanding of what was said in them, determines what was said in them—hardly a surprising fact. On the other view, the meanings of the words, e.g., 'is black', together with some (as yet unidentified) facts about their speaking—where these facts are ascertainable independent of ascertaining how the words as spoken were to be understood, *and* specificable in advance of any speaking of the words—determines how those words, on that speaking, were to be understood. For example, one may think he is able to say in advance that it is the intentions of the speaker, or that plus some other factor, *Q*, say, which together with the meanings of the words (and, perhaps, various conventions) determine in general, wherever a speaking takes place, what was thereby said. There is no particular good reason known for holding such a view, though many appear to do so anyway.

The difference between the present view of saying and what Grice appears to require might be brought out as follows. Grice certainly allows that what words mean does not fully identify what is said to be *so* in speaking them. For what is said to be *so* depends, at least, on what (times, objects) is referred to, or spoken of in *so* speaking. And that, depending as it does, on such things as the time being spoken *at*, is not in general determined by what the words spoken mean. But Grice *does* appear to take it that what words mean *does* identify something, in the saying of which that which *was* said to be *so* in a speaking *was* said to be. For it *does*, on his view, identify what was said insofar as this *is* identifiable in terms of what was to be understood as having been said, and that alone. Let us put things the other way around. Let us call a fact *nonconventional* relative to a speaking if and only if its holding does not entail, by itself, of any of the words spoken (or of what was said) that they are to be understood in any particular way. For example, the fact that a speaking occurred at 3 p.m. does not by itself, entail that any words (or anything that was said) are to be understood as referring to (e.g., something's happening at) 3 p.m., since that will depend on whether any words

are to be understood as referring to the time of the utterances. Similarly, the fact that a speaker had, say, Joop Janssen in mind when he said 'Joop is coming', does not, by itself, entail that any of his words are to be understood as referring to Joop Janssen, since that will depend on whether any of his words are to be understood as referring in such a way as to refer to the person he has in mind. (Even if it were a rule of our language that words could never do otherwise, that would be a rule about how speakings of our language are to be understood, and one we can easily conceive being otherwise.)

Grice appears to take it, then, that what is determined by the meanings of the words spoken is something such that, if we knew it, we could determine what was thereby said to be so, in terms of nonconventional facts alone (if we knew enough of them). Needless to say, this would have to be something which determined which nonconventional facts were *relevant* to fixing what was said to be so, and *how* they were so. For example, our knowledge of English tells us, according to Grice, that to know what was said to be so in a speaking of 'He is in the grip of a vice', we need to know the time of that speaking. That nonconventional fact will tell us when 'some male' was *said* to be gripped by a vice, and it will do so according to the rule: the time spoken *of* in such a speaking is identical with the time of utterance. Given what the words spoken mean, one thing, at any rate, that is *not* required is further selection between different ways in which the words, so spoken, might be to be taken. We need only determine, given the correct way of taking them, what was said to be so of what. And that, as we have seen, is what is just not so.

What is wrong with Grice's view of saying shows itself in what he takes our abilities to be for reporting what *would* be said in given words, on given meanings, in speaking about which further facts are as yet undisclosed. Suppose, for example, that I now know that, sometime in 1985, Reagan will say of someone 'He is in the grip of a vice', on the 'nasty habit' reading of 'vice'. Then, with no further information about that speaking, I can now describe what will thus be said to be so as follows: It will be said that the person Reagan will have spoken of will be, at the time of Reagan's speaking, in the 'nasty habit' sense, in the grip of a vice. In so speaking, on Grice's view, I will have characterized some particular state of affairs which Reagan, in so speaking, will have said to obtain, and which will, then, either obtain or not. As for whether I or my words have 'fully identified' what he will have said to be so, or what must be so for things to be as he will have said them to be, what I would have said in so describing Reagan's future speaking might be compared usefully with the following case. Suppose I now say that the next prime minister of Belgium will have big problems. Given our present knowledge of political facts of life, we cannot, I suppose, now identify who my words 'the next prime minister' refer to. Given an assemblage of candidates for that role, we cannot now *know* which one I have said will have big problems. After the next election, though, if not sooner, we will find out who my words *referred* to—Eddy Merckx, say—and that I

have said of him what is certainly so. We will find out, that is, that my words *do* fully identify him, given the facts, as they will emerge, that no understanding of my words is capable of changing. The only sense in which my words do not *now* identify Eddy Merckx is that we cannot now tell that that is who they refer to. Similarly, given *some* choices of possible states of affairs—Ford's being in the grip of a vice, Kissinger's being in the grip of a vice, Billy Graham's being in the grip of a vice, etc.—we cannot *now* tell which of these my description identifies, or refers to as the one Reagan will have said to obtain. As the facts about Reagan's speaking emerge, however, if he has managed to say anything in doing so, we will find out that my words, properly understood, do pick out just one of these possibilities, and in so doing, refer to just one state of affairs, as opposed to any other that might be said to obtain.

But now let us suppose—in line with what we have already seen to be so—that any of a variety of distinct personal characteristics might be spoken of in the words 'is in the grip of a vice', *on the nasty habit reading*, and any number of distinct things thereby said of someone. If my words 'that the person Reagan, etc.' do describe some particular state of affairs, then my words 'is in the grip of a vice' must speak of *one* of these characteristics. And if *Reagan* will have said anything in particular to be so in *his* speaking, then his words 'in the grip of a vice' will also speak of just one of these. But then, I will have correctly reported what *he* will have said to be so only if the one of these which my words speak of, and the one which his words speak of, happen to be the same—that is, only if both our speakings of the words happen to be understood, out of the many ways possible, as making the same contribution to what was said in speaking them. And that is something we cannot know merely by knowing which of their meanings those words, as spoken by me, and as spoken by Reagan, bear. If Reagan uses the words correctly, he will, in doing so, have ascribed to someone *some* property that might be described as 'being in the grip of a vice'. But what, precisely, it will be, and what he will have said in doing so, is at this point anybody's guess. What it will be depends on the as yet undisclosed facts about his speaking—specifically, whatever facts will determine how *his* words, as spoken, are to be understood. If, in my words, now, there is something in particular that I have described him as saying, and if that happens to be what he *will* say, that is sheerest accident.

None of this need matter very much for Grice, were it not for the fact that his view of saying systematically leads him astray when he wishes to argue for the presence of implicature. The reasons show themselves, in a general way, in what he considers evidence for implicature to be. One alleged sort of evidence for something's being implicated rather than said, for example, is what he terms 'cancellability'—the fact that a particular thing which is either said or implicated in one speaking of given words is neither said nor implicated, nor otherwise to be understood for something's being implicated rather than said. Certain aspects of what is said in given words *must* be such that what they are may vary from

one speaking of the words to another. And any feature or part of what is said which is thus variable is *ipso facto* cancellable. This is, then, not a feature of what is implicated, which, in general, distinguishes it from what is said.

Another alleged sort of evidence is what Grice calls 'nondetachability'. If, in a given speaking of words, *w*, something, *x*, would be to be understood, and if it is not 'possible to find another way of saying the same thing', in the same circumstances, but in other words, such that *x* is not part of what is to be understood, then this is at least some evidence that *x* is a 'generalized conversational implicature'. All this is a bit vague, depending, as it does, on what one counts as saying the same thing, or being the same circumstances. But at least, here again, whatever goes for conversational implicatures goes equally well for the *variable* aspects of what is to be said in given words. For any such variable aspect, and for any given speaking, the selection from out of the range of possibilities of things to be said in the words in question of that particular possibility which *does* correctly represent what was in fact said must be affected by some facts or other about the circumstances of the particular speaking in question. Now suppose that those facts, whatever they are, remain the same, while we substitute for the words spoken other equivalent, or nearly equivalent ones. It will not be surprising if nearly equivalent words represent a highly overlapping class of possible things to be said. But then, it should also not be surprising if under these circumstances, the same facts frequently make the same selections.

Finally, according to Grice, establishing that something, *x*, is conversationally implicated rather than said in a particular speaking requires constructing an explanation of how it could be worked out, in terms of facts about that speaker, and in particular, in terms of the putative observance or nonobservance of what he terms 'conversational maxims' that *x* would be part of what has to be understood by or in those words as thus spoken. But again, the availability of such explanations would provide absolutely no evidence in favour of putting something on one side or another of a 'what is said'–'what is implicated' dichotomy. What such an explanation would show, if available, for a particular aspect of what there is to be understood in or by a speaking—e.g., that in saying the kettle to be coloured black, it was or was not being said that it was made of a black substance, or that it was merely black on the surface—is that that aspect needn't be so to be understood *in virtue of the 'conventional' meanings of the words spoken*—which might reasonably lead one to the hypothesis that it is not what the words mean, e.g., in English, which makes them, on that speaking, so to be understood. But this, of course, is a feature which is *shared* by what is conversationally implicated (whatever that might be) and those features of *what is said* which may vary from one speaking to another of given words. Once one recognizes that what, e.g., English words mean may *not* choose between a wide variety of distinct things there are to be said, any of which may therefore be said in speaking them, such an explanation, by itself, is no longer any evidence that something is not, strictly speaking, a part of what is said. For it may

simply explain why that which was *said* was, in fact, to be taken as that which was said.

To see how the enterprise of finding implicatures is truly touched by such considerations, it is best to descend from the above level of generality and look at at least a specimen of actual practice. Grice's misreasonings about what is said and what is implicated, prompted by his misperception of the connection between what words *mean* and what is to be said in them, are illustrated by remarks he makes on the English indefinite article. On this topic, he first notes, correctly, that sometimes (e.g., for some speaking of 'Sam entered a (someone's) house and found a turtle on the coffee table'), one would be surprised to discover that the house (or the turtle) which the speaker had in mind, or which Sam was thus involved with, was Sam's own. Other times (e.g., for some speakings of 'Max broke a finger') one would be equally surprised to find that the finger in question was *not* one of Max's own. Yet other times (e.g., for some speakings of 'Sam spilled a beer'), nothing is in any way suggested about whether the beer in question was Sam's own or not. Grice further notes, also correctly, 'I am inclined to think that one would not lend a sympathetic ear to a philosopher who suggested that there are three senses of the expression *an X*' (Grice 1975, p. 56).

Both of these observations would be, not merely true, but also interesting, were it the case that variations in what is said in a given form of words must correspond to a variety in the senses which those words may bear in their home language (e.g., English). For, given that there are no three senses of the indefinite article in English, we could then conclude that whatever variation is exhibited by the speakings alluded to above, it could not be a variation in what was said, or at least not one induced by the presence of the indefinite article.<sup>12</sup> And in that case, we would be inclined to lend a *very* sympathetic ear to the thesis that the observable variation is actually a variation in what is being conversationally implicated in various speakings involving indefinite articles—just the conclusion which Grice suggests we *should* draw.

The only difficulty with this line of thought is that it is simply *untrue* that variation in what is said in different speakings of given words need signal variation in the meanings (e.g., in English) which those words bear on the various speakings. Suppose, then, that someone wanted to claim that in various speakings of 'Max broke a finger', 'Clothilde lost a purse', and so on, different things; strictly speaking, were said, some of which would be true if, e.g., Clothilde lost her own purse, but not if she lost someone else's, for others of which the opposite would be the case; some of which would be true if, e.g., Max broke someone else's finger, but not if he broke his own (e.g., speakings reporting on Max's daily chores as Mafia enforcer), and for others of which *this* particular distinction is not relevant, but others may be. A sensible account of what is to

<sup>12</sup> The same points, of course, hold, as indicated by the examples, for other devices for expressing existential quantification—e.g., 'someone'.

be said in speaking an indefinite article on which such claims turn out to be true might look like this (I claim merely that it is sensible—not necessarily that it is true): Any speaking of an indefinite article is to be understood, in part, in terms of a more or less restrictive and more or less definitively determined class or sort of individual, such that it is only an individual of that sort, or within that class, which is being spoken of *via* the noun phrase which the article introduces, and thus only such an individual which can thereby contribute to making what was said in the speaking containing that of the article true, or to making it false. If one likes, one could think of a speaking of ‘a(n)’ as introducing an implicit reference to such a class. One could even think of speakings of the form ‘a(n) X’ as modelled on expressions of the form ‘the tree’, or ‘Jones’, or ‘the first person to fry whitebait in a djao lei’, as carrying with them conditions of one sort or another for being the right *sort* of X, just as in the latter cases, such referential speakings carry with them conditions for being the *right* tree, or the right Jones, or the right time or occasion for being the first person to fry whitebait in a djao lei—assuming, that is, that anything *could* be the right tree, or Jones, or etc.<sup>13</sup>

However one thinks of it, some specification of such matters will, on the hypothesis, be part of a correct account of what was *said*. Some such specifications will be such that, e.g., a purse other than one of Clothilde’s would fall outside of the relevant class, so that her having lost that could not make what was said true. Other times, such a purse would fall inside the relevant class. And other times, any of countless other distinctions might be marked through such a specification. What is important for present purposes, however, is that none of the variation in what is said described by specifications of this sort corresponds to any variation in what a phrase of the form, ‘a(n) X’, is to be taken as *meaning* (in English) on the relevant speakings. On the contrary, we explicitly assume that there is no more than one meaning of the English indefinite article throughout.<sup>14</sup> Thus, Grice’s second observation in no way counts against such a hypothesis, nor in

<sup>13</sup> This would, in turn, of course, raise a whole host of issues about whether the conditions are to be represented in terms of ‘descriptive backings’ to be understood for each such speaking, or causal relations between speakings and candidate objects, or something else. Then there would be further issues about whether ‘condition’ in the above ought to be taken for something as strong as a necessary and sufficient condition, or as something essentially weaker. (My own view is that it is some notion other than necessary and sufficient condition which is appropriate in such matters. I have expressed that view, in the case of reference, in Travis (1981*a* and 1981*b*.) At any rate, these are issues we would have to face and sort out once we ceased to dismiss the phenomena in a Gricean manner, and came seriously to work out hypotheses of the above kind. A djao lei, by the way, is a particular kind of Chinese kitchen implement used in deep frying or boiling.

<sup>14</sup> Or, perhaps preferably, one contribution which the definite article makes to the meanings of phrases in which it occurs. This does not rule out that there might be *other* ambiguities in constructions involving the indefinite article—e.g., to cite a famous case, two *readings* of ‘Sam wanted to buy a house’, and not merely a variety of ways in which such speakings might be to be understood. The alleged ambiguity there *might* be handled as a special case of the variation described in the above hypothesis. But whether it should be or not is a difficult issue, into which I will not attempt to enter here.

favour of any rival one on which the observable variation from speaking to speaking is variation in the conversational implicatures present.

But what would be the *evidence* for the correctness of a hypothesis like the one just sketched? That depends in part, of course, on the rival hypotheses in the field at a given time. But it need not be *quite* as direct as that envisaged in Grice's first observation. Consider a case like the following: Clothilde returns one evening from the office and says to Sam, 'The police were around today. A wallet was missing.' Now suppose that no wallet was missing at the office that day. Would we (strictly speaking) count what Clothilde *said* as true? Surely, somewhere in the world, someone's wallet was missing (or it was a very off-day for pickpockets). But that is surely not enough to make what she said *true*. If it were, it would be all too easy (and correspondingly useless) to speak the truth *via* existential quantification. On the contrary, despite empty pockets in Bangkok, Capetown, and Pyonyang, we would have no hesitation in terming what Clothilde said false. Apparently, then, there is a restriction of the kind envisaged in the thesis as to what would make what she said true. Looking for the nature and the source of the restriction, one comes up with a bounty of cases which *contrast* with the above. Suppose, for example, that Clothilde and Sam were the previous day in Bangkok, where they were detained and interrogated for allegedly suspicious behaviour in the lobby of a posh hotel, and eventually released, provisionally, with the admonition that their local police would be in touch if anything turned out to be missing at the hotel. Then, Clothilde may well have said in her words what it would take a wallet missing in Bangkok since the previous day to make true. Or suppose that Clothilde is a super sleuth. Police from all over the world consult her, through their local counterparts, whenever anything turns out to be missing in a truly baffling and important way. Then the relevant place for a wallet to be missing, with respect to what she may well have said, might be virtually anywhere, though there will presumably be other restrictions on the sort of wallet it may be, or its connections with Clothilde's life. Though the cases by which we thus build contrasting n-tuples with respect to the indefinite article may sometimes appear a bit strange, they serve our purpose of illustrating the various things there are to be said, in some words or other, by a given speaker to a given audience, at a given time. Since none of these things is tied to the *meaning* of the indefinite article in a way that any other is not, the conclusion (pending strong arguments to the contrary) is that all of these are distinct things to be said in the way they appear to be sayable—*via* speakings of the indefinite article—and that, consequently, a thesis somewhat like the one sketched above—as opposed to one merely about the presence of implicature—must be correct.

To summarize, to arrive at a proper account of the indefinite article (and *mutatis mutandis* of any other linguistic device), first look carefully at the distinctions there are to be drawn between one thing and another which there is to be said, and which is the correct understanding of some speaking or other



of words employing the relevant device. Determine what (sorts of) feature one would need to mention in order to distinguish correctly between one such way of understanding and another. One may then pose a variety of questions. First, to what extent is it reasonable to suppose that these features would have to be mentioned explicitly in the words spoken if that which there was to be said was actually to be said, and not just intimated or suggested, or etc. Is it reasonable to suppose, for example, that Clothilde would have had to have said in so many words that a wallet was lost in such a way as to cause some police agency to get in touch with her in her capacity as sleuth for her to have *said* what only a wallet lost in that way would make true? As a subquestion, is it reasonable to suppose that a speaker would be in a position to *know* how to state correctly the distinguishing features of one of these things there are to be said—e.g., the correct conditions on a lost wallet's being relevant to the truth of one such thing or another? Second, are the distinguishing features involved such that it is reasonable to take those attaching to one particular such thing to be said—as opposed to others that there are to be—as those features 'really' determined by, e.g., what the words, 'someone lost a wallet' *really* mean, in a way in which those features attaching to other things to be said would not be? Is there, for example, some restriction on the contextually relevant class of lost wallets which distinguishes itself as the one which the English words 'lost a wallet' really mean? These are not the only possible fruitful lines of thought. One totally unfruitful line, however, as far as determining what is said is concerned, is to ask whether the presence, in a given case, of one set of distinguishing features or another is predictable or explainable from some other features of a speaking, distinct from what the words spoken mean. For what that leaves entirely open is whether *what* gets thus explained is the saying of what was said, or the suggesting, intimating, implying, or etc. of something else.

### On What Saying Is

There are two ways one might count things as said or not. One is to rely on our intuitions, asking such questions as 'Would one say that the speaker (actually) *said* X?', or 'Was X *said* or not in that speaking of those words?' Where answers are clear enough, we can, accordingly, say either that X was said or that it wasn't. The other way is to count X as a part of what was said in given words just in case mentioning it is necessarily part of adequately specifying what *was* said in those words—where this means (i) adequate for distinguishing what was said from whatever else might have been in speaking those words, in speaking that language, but was not, and (ii) adequate for specifying what was said to be *so* in the words—or, if we name what was said to be so Y, adequate for saying what it is that must *be* so for Y to obtain. There is nothing wrong with either way of regarding saying, but the two ways are liable to yield rather different results,

and focusing on the first is likely to mislead when it comes to separating off that which was implicated from that which was said.

Grice, for example, is impressed by the fact that if he says, of some A, 'He is an Englishman; he is, therefore, brave', then he would not say that he had *said* of A that it followed from his being an Englishman that he was brave. Suppose Grice is right not to say this (i.e., that it isn't so). Grice apparently takes this to mean that, while he had 'indicated, and so implicated' that this follows, 'I do not want to say that my utterance of this sentence would be, *strictly speaking*, false should the consequence in question fail to hold' (Grice 1975, pp. 44–5). *Perhaps* Grice is correct in this too. On the other hand, his first claim *may* be correct even though his second is not—even though, that is, what he would have said in the above words *is* strictly speaking false. For it may be that what was said in those words was that A was an Englishman and he was brave, where *this was to be understood as following*, perhaps in some more or less specified way, from his being an Englishman. And it may be that this was so to be understood that what must be so for what Grice *said* to be so to be, is that A be an Englishman, and he be brave, where his being brave follows in an appropriate way from his being an Englishman. So the latter's following from the former may be *part* of what was said, in the above sense, and part of what must be so, for what was said to be, strictly speaking, true, even though it was not *said that* the latter followed from the former.

One does not, as a rule, *say*, in speaking them, how his words are to be taken; he simply speaks them so that they *are* so to be taken. Suppose Reagan to say 'Gerald is in the grip of a vice'. He did not, we may suppose *fail* to say which Gerald was being spoken of. But nor is there anything he said in which it was said *that* the Gerald being spoken of was so-and-so. Suppose it was Ford. He did not then say *that* the Gerald being spoken of was Ford (though, since it was, there is something else he said about Ford). In nothing that he said, in fact, was anything said about how his words were to be taken. Nor was it *implicated*, e.g., that the Gerald being spoken of was Ford: not conventionally, since nothing about the meaning of the name 'Gerald', insofar as it has one, makes it so to be taken that it means Ford—not even as spoken by Reagan; nor conversationally, since nothing was said, aside from that in which 'Gerald' did mean Ford, on the basis of which that fact might have been worked out. 'Gerald', as spoken by Reagan, simply *was* so to be understood that it meant Ford—neither by this having been said to be so, nor by its having been implicated. (And similarly for other aspects of what his words were to be understood to say—e.g., that 'in the grip of a vice' was to be understood as attributing one property rather than another which it might have.)

What Reagan said in virtue of this last fact was that a certain Gerald was (taking this in an appropriate way) 'in the grip of a vice', where the Gerald this was, was Ford. Now, what must be *so* for what Reagan said to be true? Surely not just *any* Gerald's being 'in the grip of a vice' will do. The Gerald thus gripped

must be the right one, where the right one, in this case, is Ford. But why is this so? Not because there is anything which Reagan said in which it was *said that* Ford was the right one—i.e., the Gerald so gripped—but rather because it was said that a certain Gerald *was* so gripped, where that Gerald was to be understood to be Ford. And this aspect of what there was to be understood—that the Gerald in question was Ford—was to be understood in such a way as to contribute to a specification of what must be so for what Reagan said to be so to be—and hence, for what he said to be *strictly speaking* true. It is thus *part* of what was said, in the sense outlined above.

Its being to be understood that someone's being brave obtains in a particular way (e.g., as a consequence of his being British)—where this *is* to be understood—*may* contribute to fixing what was said to be *so* in given words (e.g., 'He is, therefore brave'), in much the same way that its being to be understood that the Gerald in the grip of a vice was Ford contributes to what was said to be so in Reagan's. In particular, it *may* contribute in such a way that what is thereby to be understood (according to those words) as being so—that his being brave obtains in a particular way—must *be* so, for what was said in those words to be true. This would certainly account for the (my) strong intuition that if, e.g., Englishmen are notorious cowards, and their upbringing notoriously apt for keeping them that way, then one might respond correctly to Grice by saying, 'It's not true that he's an Englishman and therefore he's brave. He's brave alright, but *despite* being an Englishman.' At least such intuitions are not to be dismissed (e.g., as not *really* about what was said) *merely* on the grounds that there was nothing Grice said in speaking those words in which it was said *that* A's being brave followed from his being an Englishman.

None of this is to assert that it is correct that the bravery being a consequence of being English is, in fact, part of what must be so for things to be as said in Grice's words. To establish such a hypothesis one would, *inter alia*, need to note the similarities and differences between what is to be said in Grice's form of words and what is to be said in other closely related forms in which roughly the same points might be made—e.g., 'He is an Englishman, from which (fact) it follows that he is brave'; 'He is an Englishman, and, as a consequence of that, he is brave'; or 'He is an Englishman, and it follows from that that he is brave'. Whether any of these and Grice's words differ enough from each other that it takes different things being so to make what is said in each true is, I think, less than obvious. But actually *doing* the work from which Grice wishes to excuse us is not our present task. Here I will merely note that situations like that sketched above—that is, situations in which, while it wasn't exactly said *that* X, X's being so is part of what must be so for that which was said to be so to obtain—*must* arise reasonably often.

Consider a speaking of the words 'Max refused to ask Clothilde in', in which it is said that, on a particular occasion, Max refused to ask Clothilde in. What must be so for what was thus said to be so to be? For one thing, it certainly must be so

that, on that occasion, Max did not invite Clothilde in. Beyond this, normally, a number of other things would be at least suggested—e.g., that Max was at least aware of Clothilde's existence, that he had some kind of an opportunity to ask her in. In fact, verbs from Grice's favoured class, such as 'suggest' and, particularly, 'imply' seem to fit rather naturally for the ways in which such points were conveyed. But now let us turn to matters of the truth of what was said. Must any of these things actually be so, for what was said, strictly speaking, in those words, to be true? Consider the question of opportunity. Nowhere in what was said was it said *that* Max had the opportunity to invite Clothilde in. By this test, then, that is not part of what, strictly speaking, was said. We might, then, relegate this suggestion to the realm of implicature (perhaps of a conventional sort), with the consequence that the suggestion, strictly speaking, may be false, while what was actually said is true, and hence that Max's having had the opportunity is not part of what must be so for what was said to be so to be, nor for what was said to be true.

If we apply this test systematically, we will discover that none of the above suggestions, nor anything like them, is strictly speaking part of what is said to be so in speaking the above form of words. And this begins to indicate why the test cannot be a valid one (aside from the fact, e.g., that if Max is a typical Winnipegger, while Clothilde has spent her entire life cloistered in a nunnery in Nepal, we wouldn't say it was *true* that Max refused to ask her in). It cannot be, that is, that failing the 'it was said that' test has, as a general consequence that something was not part of what was said about Max—e.g., things said in speaking the words 'Max did not ask Clothilde in', 'Max did not manage to ask Clothilde in', 'Max forgot to ask Clothilde in', or 'Max failed at asking Clothilde in'. One thing that must be so for things to be as said in any of these is, again, that Max did not ask Clothilde in. Any suggestions over and above this, however, are more than likely to succumb to the same fate as those in the case of 'refuse', with the consequence that what is said in *any* of these will be, strictly speaking, true just in case any of the others is. In the case of 'forgot', for example, there is a suggestion, at least, that Max once had it in mind to ask Clothilde in; in the case of 'failed at', a suggestion that some sort of attempt was made. But nowhere in the above remarks was it *said that* such things were so. So, on the above test, such things cannot be part of what was said to be so either.

Such results would not show, of course, that, e.g., the English words 'failed at asking' are synonymous with the English words 'refused to ask', nor that either mean the same as the words 'did not ask'. For each of these expressions may still differ, at least, in what conventional implicatures it carries. The trouble comes with the *contrast* Grice wishes to draw between implicature and matters affecting the truth or falsity of what was *said*. For, what the test *would* show to be so, which is not, is that what was said to be *so*, e.g., in saying, 'Max refused to ask Clothilde in', does not differ from what was said to be *so* in saying, e.g., 'Max did not ask Clothilde in', or at least not in those ways in which it obviously does. And that is why the test is not a good one.

Suppose Max could not have reached Clothilde, and anyway, had never heard of her. Then it is not *true* to say that he refused to ask her in, though if we ask of what contradicts our supposition (e.g., that Clothilde was reachable) whether it would be *said* in saying so, the answer, according to our first criterion for being said, is ‘No’. The conclusion is modest: The tests ‘Was X *said* in speaking w?’, or ‘Was it said in w that X?’, cannot, where they yield a negative result, be taken to show that X’s obtaining is not part of that state of affairs which was said to be so in w, much less that it is not part of what must be so for what was said, in w, to obtain, to do so. Perhaps Grice has other reasons for such claims in particular cases. If so, it would be interesting to see what they are.

### A-Theses

Conversational implicature was meant to have import for a certain type of philosophical thesis, termed an ‘A-thesis’. Though Grice identifies the type by one particular feature, typically, at least when intelligently formulated, a thesis has that particular feature in virtue of having a number of others. And in virtue of the others, the intended import of conversational implicature does not exist. To say that is not to endorse any particular A-thesis, though I confess to sympathy with a considerable number of them. But the way to see the point is to set out at least the skeleton of a particular A-thesis in some detail. For that purpose, I choose a sample thesis about actions done voluntarily.

At the heart of an A-thesis is the view that it is a mistake to think that one could survey the relevant objects—doings in general, in this case—and make an inventory of those done voluntarily and those not so done—that it is a mistake to think that doings as such dichotomize in such a way that some such inventory would be the right one, doings falling on the one side being those that really *are* truly describable as ‘done voluntarily’, and doings on the other side being those that are not. The reason that this is thought a mistake is that the properties which doings have do not, by themselves, make it correct to describe the doings in the one way or the other. And the reason for that is that there is no such state of affairs (which, for a given doing, might *hold* or not) as being what something would be said to be in describing it as done voluntarily.

The reason for this, in turn, is that there *is* no thing in particular that something *would* be said to be, in describing it as done voluntarily, where that something is determined merely by the fact that the description is as being done voluntarily. At this point, we come to the central linguistic thesis that needs to be made out in establishing the central intuition behind our A-thesis. And that is this: Let ‘D’ be a way of referring to a doing, and D a doing which it might refer to. Then there are *many* distinct things to be said in saying ‘D was done voluntarily’, where ‘D’ refers to D, ‘done voluntarily’ means *done voluntarily*, and where some of these things to be said will be true, and other of them will

be false. Thus, some descriptions of D as done voluntarily will say what is true; others will not. Which is to say, a doing, D, may be done voluntarily, given one of the things that being done voluntarily may correctly be regarded as coming to, and may not be done voluntarily given another one of these things. But in general, it will neither be 'done voluntarily', nor 'not done voluntarily' *tout court*, since there is *nothing* that this comes to.

The above sort of linguistic thesis, as well as its consequences, are by now familiar. The point is that *all* of the distinct things to be said in the above-mentioned speakings are things to be said *in describing D as being done voluntarily*. None speaks merely of a sort of *soi-disant* or 'courtesy' voluntariness. Hence, none of these things to be said contrasts with the others in revealing, as they do not, what being done voluntarily 'really' comes to. But then, choose any subset of the set of doings. Given the various things *thus* to be said in speaking of its members, there will be no good reason to take *that* set as 'the class of actions done voluntarily'.

Of course, where some speakings of the words 'D was done voluntarily' occur, some of the things-to-be-said in some speaking of those words are very likely to be appropriate things to say, and others are very likely *not* to be. In any case, for any such speaking where something in particular, in fact, *was* said, *some* facts about the speaking will select just one of these things-to-be-said as that which actually was said. And one might speculate that, where this is so, that will generally, if not invariably, be an appropriate thing, rather than an inappropriate thing to say on that occasion. At any rate, a specification of what *was* said, if it is individuating, must be in terms of some features of what there is to be understood in properly understanding the words which go beyond anything determined by what the words conventionally mean. Such a specification may itself mention features of the speaking in question, or it may not, depending *inter alia*, on the A-thesis in question, and what it has to say about what possible specifications are like. If it does mention features of a speaking, those features will, in general, be distinct from whatever facts about the speaking make it correctly taken as saying one thing or another, and hence, which make that specification the right one.

We now come to a subsidiary and independent part of our skeletal A-thesis, and it is just this part which catches Grice's eye. To what has been said so far, we now add the following: It is possible for a speaking of the above words, which manages to refer to a particular doing, D, and in which 'voluntarily' means *voluntarily*, yet to be done in such a way that *no* facts about it select any of the possible things to be said in the words as that which actually was said. In this case, *nothing* was said, or at least, nothing was said to be so, and hence nothing said which might either be true or be false. Note that *this* part of the A-thesis, plausible as it is, *might* be false while everything else said so far in stating the A-thesis was true. Maybe it is simply *impossible* to speak certain words and thereby say nothing—though the evidence from simpler cases is against this. Suppose, for example, that someone says 'Gerald found a solution to the problem'. As we have

seen already, any of many things might be what was thereby said. One variable distinguishing one such thing from another is what, on that speaking, 'Gerald' meant. On some speakings, it will be so to be taken that it meant Ford; on others, so to be taken that it meant (some) Gerald Cohen. But suppose I just say, now, for no reason in particular, 'Gerald found a solution to the problem'. No facts about this speaking determine 'Gerald' (or 'the problem', for that matter) as thus spoken, to mean one thing, as opposed to any other which some such occurrence of the words *might* mean, and certainly not such that those words, as thus spoken, mean any one (thing). And, as an obvious result, there is nothing that I will have thus said to be so at all. Perhaps the case is different with 'done voluntarily'. But even if it is impossible to say *nothing* (either true or false) of D in saying 'D was done voluntarily', it may still be possible to say any of many distinct things, in all of which D is described as done voluntarily. And that is really all that an A-thesis, at heart, requires.

Let us now add a further subsidiary, and yet stronger element, to our A-thesis, namely, the following claim: Pick any speaking of 'D was done voluntarily', in which some particular thing was said of D, in describing it as done voluntarily. Then there will be at least some doing, E, such that, if we had substituted 'E' for 'D' in the above speaking, so that 'E' referred to E, while keeping as many other facts as possible about the circumstances of the speaking constant, *nothing* in particular will have been said of E. In other words, if we think of our original speaking as done on an occasion where we set out to classify *every* doing either as done voluntarily, or as not, there will be at least some actions of which, even given what is fixed about what is to count *on that occasion* as done voluntarily or not, nothing sensible will have been said in calling them done voluntarily or not, and hence which, even given our current interest in classification, cannot sensibly be placed on one side or the other of our projected dichotomy. It should be clear how this thesis is stronger than, and independent of, the heart of our A-thesis, which is simply the intuition that an action cannot count as voluntary or not *independent* of something in particular being meant by this, where any of many distinct things, correctly, might be.

A-theses, being *philosophical* theses, there are, initially, two sources of evidence for them. One source is simple observation of linguistic phenomena: Suppose John handed over the roll of bills without having his fingers pried loose from it, but under the threat of death. Then we can construct cases where something *true* was said, and cases where something *false* was said in saying, 'John handed over the bills voluntarily'. The other source is the existence of a philosophic problem. In brief, for a few thousand years, philosophers have been unsuccessful in drawing a dichotomy, in a convincing way, between what is voluntary and what is not, on which it does not turn out either that everything (or far too much) is voluntary, or that everything, or far too much, is not. Given the unacceptability of either outcome, one might become inclined after a while to take this failure as evidence that no such dichotomy is to be drawn.

One might, of course, challenge the *central* part of an A-thesis, denying, e.g., that a variety of distinct things really are to be said in classifying a given action, A, as voluntary. And one might try invoking the notion of conversational implicature in doing so, explaining the variation in what one would understand from case to case in the data uncovered by the A-theorist as really variation in what is conversationally implicated, rather than in what is said. But here one runs up against a problem exactly parallel to that which faced a proponent of conversational implicature in the case of 'is black'. Suppose the A-theorist to have assembled, plausibly, contrasting n-tuples of reasonably impressive size and weight—e.g., a good collection of potential speakings of 'is done voluntarily', speaking of a given action, in each of which a different thing is at least apparently said—for each of which, that is, there are distinct conditions under which we would be inclined (albeit perhaps naively) to say, 'That's true' ('false'), 'what Harry *said* is right—he *did* do it voluntarily (but what he didn't say is more important)', etc. Then, minimally, the A-theorist has unearthed a good variety of distinct things which there are to be said—whether or not they are to be said in those words, or in saying that D was done voluntarily.<sup>15</sup>

Now, before anything can be conversationally implicated, *something* must have been said in the words concerned, in each of the A-theorist's cases. The question is, what? If the A-theorist is to be wrong, then it must be the same thing which was said in all of them, provided that, in each, the words were used as meaning what they *do* mean in English. But then, which of the A-theorist's candidates for this post should it be? Or should it be something else altogether? If the implicature theorist is to make his case, then there must be some principled way of deciding this, which means that there must be a fact of the matter of the sort he requires—i.e., what the right choice is must be determined either by the meanings of the words 'voluntary' and 'voluntarily', or by our concept of being voluntary, or what we recognize it to be to be done voluntarily, or something of this ilk. But, as in the case of 'is black', if the A-theorist has built his case carefully, these are all sources from which we are likely to get little guidance. The issue is not apt to be whether 'voluntary' meant *voluntary* or something else, but rather what counts as being *that* (i.e., that which 'voluntary' means)—and that is precisely what is liable to vary from instance to instance of classifying doings as those to which our notion of voluntary applies, and those to which it does not.

This, of course, is a sketch of an issue, and not an arguing of it. Though arguing it is not the present point, clarity may be served if there are at least some actual sample arguments on hand. To generate some, let us take up an issue yet closer to Grice's heart—that of the proper understanding of (some) sentential connectives, that is, words which function grammatically to connect sentences, such as the English 'or', 'and', 'but', 'if... then', 'despite which',

<sup>15</sup> A useful principle: what was implicated *might* have been said, in some words or other.



‘since’, ‘because’, ‘before’, ‘after’, and ‘where’. For the first four, at least, Grice wishes to defend the thesis that the contribution made to what is said in the speaking of words in which they occur is purely truth-functional, in the following sense at least: In general, where such a connective occurs in a speaking of given words, whether the whole of what is thus said is true depends only, in one way or another, on whether what is said in the words the connective connects is true. For example, if one speaks words ‘A and B’, in which ‘and’ meant *and*, and something was said, then what was said is true just in case what was said in speaking ‘A’ is true, and what was said in speaking ‘B’ is true. Similarly, what was said in speaking ‘If A, then B’ will be true if what was said in speaking ‘A’ is false, and also if what was said in speaking ‘B’ is true. I will focus attention in what follows exclusively on the thesis for ‘if. . . then’.

As Grice recognizes, such a thesis runs immediately against an apparently great problem. Suppose you are eyeing yet another plate of dim sum, and, as the cart approaches, I say, ‘If you eat that, (then) you’ll get sick’. Suppose you later discover that all I had in mind was that, since you had already been infected with the influenza virus, you certainly would get sick—and the consequent of my conditional thus being true, that the whole of my conditional was. You would, as Grice would put it, at least be highly surprised. Further, suppose you *did* later contract influenza, where this was clearly unrelated to your eating dim sum. One would be at least highly reluctant to allow that that made what I said true. Or, if I tell you ‘I told you so’, or ‘I said that would happen’, you might feel rather justified in saying, ‘No, you didn’t’. So there is a suggestion, *minimally*, that something else was said besides (or other than) what is represented as said on the truth-functional account—something which cannot be true merely in virtue of the appropriate truth values of the connected parts. Roughly, there seems to have been something said about some connection between your eating that last plate of dim sum and your getting sick—or at least this appears to be part of what was said.

Grice’s way with such appearances is to relegate them to the realm of conversational implicature. As he would have it, it was indeed *implicated* that eating the dim sum had something to do with getting sick. You do thus have every reason to be surprised and even annoyed when you find out what was really in my mind. Why ever did I put that point in that rather bizarre way? But though—or rather, because—something like this was implicated, it was not part of what was said, or part of what is required for the *truth* of what was. Note that this way with words *does* have the effect of dismissing at least half of our intuitions in cases like the above as incorrect, even if it is successful in accounting for the other half. Perhaps we have every right to be surprised at what was in the speaker’s mind. But any reluctance we have to allow that what was *said* is true is misplaced. In the above pictured circumstances, on this analysis, what was *said* was true, I did say to you that what would happen would, etc. And what we say in denying such things—in saying, e.g., ‘Well, anyway; what you said

in the restaurant wasn't true', or 'No, you didn't say that', is, *strictly speaking*, false.

That, e.g., a connection between eating and getting sick is implicated rather than stated (or part of what is) in a speaking of the above, is a significant hypothesis, which one might well want to check. In doing so, one once again runs against the general problem that what is (conversationally) implicated in given words depends on what is said in them, and what might be worked out on the basis of *that*. So we might raise the question, what is, or is to be said in saying 'If A, then B'? And this, as experience has shown, raises the further question whether there is just one thing, or many, to be said in so speaking, where what is said in saying 'A' and what is said in saying 'B' remain fixed. Grice, of course, offers a hypothesis about this: namely, that there is just one thing to be said in so speaking, and that what it is, is what is specified by the truth-functional account. But, for given 'A' and 'B', there is a wide array of other possibilities—things that are to be said in some words or other, and *appear* to be said in one or another speaking of the words 'If A, then B'. So, taking Grice's suggestion as one possibility, the question is what (if anything) would show it (or some other) to be the right one—that is, that which was *really* said throughout the range of observable variation in the various speakings of those words which there are to be made. What is there, e.g., about the meanings of the words, 'if. . . then' which would make Grice's choice (or some other) the right one?

The evidence is that there is no such thing as 'that which would be said in saying "If A, then B"', where what is said in saying 'A', and what is said in saying 'B' is fixed, but that, on the contrary, a speaking of the words, 'if. . . then', may make any number of distinct contribution to what is *said* in, *inter alia*, so doing. The variety of things which there are to be said, and which appear sayable by this means is illustrated by the following cases.

- a. Sam and Max are driving around, mildly lost, but in a clear general area and direction, when Max notices a sign saying 'Bredaseweg'. Sam then says, 'Ah! If this is the Bredaseweg, then we'll eventually come to Breda.' That is what happens, so, we may suppose, what Sam said is true. Imagine the following objection: Maybe 'Bredaseweg' is just a name. Maybe it ends in a field somewhere, or actually goes to Den Bosch. To this, Sam might reply, e.g., 'No. They don't use names that way in the Netherlands', or, simply, 'No, I know this Bredaseweg'. Assuming either answer is correct, the objection, whatever it shows, does not show that what Sam said was false. (Suppose that, as they were about to come to Breda, their car drove off the road and into a tree, with the consequence that they never got there. Would *this* show what is so said false? I do not think so.)
- b. Pol and Suske, two Belgians, are driving for the first time in the Netherlands. Coming to one strange village after the other, they have managed to get hopelessly lost. Pol finally spots a sign which reads 'Bredaseweg' and says,

triumphantly, ‘Ah! If this is the Bredaseweg, then we’ll eventually come to Breda.’ But Suske points out correctly,<sup>16</sup> ‘Holland isn’t like Belgium. They just name their streets anything. “Bredaseweg” doesn’t mean that it goes to Breda. It might just end up in a field, or go to Den Bosch.’ Pol replies, crestfallen (and, I think, correctly), ‘Oh. Then it’s not *true* that if this is the Bredaseweg, then we’ll come to Breda. Better ask directions.’ Suppose that, by mere coincidence, this particular Bredaseweg *does* happen to lead to Breda, and that by following it, that is where Pol and Suske arrive. That is not likely to make Pol alter his judgement of his original statement. Nor, I think, should it.

- c. Sam and Max are once again out driving, but this time with special purpose. Max is afraid to drive anywhere by himself, and with some reason: Whenever he has tried (numerous times), he has always gotten hopelessly lost. Sam is now giving him point by point instruction in how not to lose one’s way. At one point, he says, ‘Look—there’s a sign that says “Bredaseweg”. Now if this is the Bredaseweg, then we’ll eventually come to Breda.’ To this, Max replies, ‘that’s not true’. (Or ‘It’s not *true* that if this is the Bredaseweg, then we’ll come to Breda.’) ‘We could be going the wrong way, how am I supposed to find *that* out?’ Here, *my* sympathies, at any rate, lie with Max—meaning that he seems to me to have spoken truly.
- d. Max and Sam are trying to reach Breda, but running dangerously low on fuel. It is late at night, there are no service stations, and whether they will reach Breda or not is a dicey matter. Though they are not quite sure where they are, Max thinks he knows the car and the local geography pretty well. So he says, ‘Drive slowly. If this is the Bredaseweg, then we’ll eventually come to Breda.’ But Max is wrong. The Bredaseweg is longer than he thought, and the car more fuel-hungry. There is not enough fuel left to make it to Breda on the Bredaseweg. Fortunately, Max and Sam are not *on* the Bredaseweg, but on another, and much shorter route to Breda. Hence, they do make it to Breda. Despite that, we have a strong inclination, I think, to say that what Max said is not true.

There is, then, something to be said, in some words or other, which is false if Bredasewegs do not generally lead to Breda, and something else which is true, provided only that *this* one does; something that is false if the car ends up against a tree, and something which is not; something which is false if it is true that the travellers may be going in the wrong direction, something that is false if there is not enough fuel to reach Breda on the Bredaseweg, even though Breda is eventually reached, and other things which might be true despite such difficulties, to merely scratch a surface. And all these things, among others,

<sup>16</sup> Trust me for such facts, which I make up freely as I go along.

*appear* sayable in given words, 'If A, then B', while what is said in what the 'if. . . then' connects remains constant. Pending some powerful reason for taking things to be *other* than they appear, such examples, made out in a convincing way, provide strong evidence that there are many distinct things to be said in speaking the words 'If A, then B', where what is said in speaking 'A' and 'B' remain fixed.

That there are many distinct contributions to what is said, any of which may be made in some speaking of the words, 'if. . . then', and that there are distinct things thus to be said, depending on what the contribution is—*if* this is so—has the following significance: If it is so, then what is said in a given speaking of words 'If A, then B', given what is said in thereby speaking 'A' and 'B', *cannot* be determined by the meaning, in English, of 'if. . . then' alone, but must also be distinguished (and specifiable) in terms of extra distinguishing features of what is said—features which may *vary* from one such speaking to another, and which *will* vary wherever, in any two such speakings, different things are said.

One might put it this way: the *meanings* of the words, 'if. . . then' plus the *thoughts* expressed in a particular speaking of 'A' and 'B' do not determine that any particular thought was expressed in what was said in a speaking of the words 'If A, then B', but leave it open that any of many distinct things might have been, depending on what other distinguishing features were to be understood as individuating what was said in that speaking and distinguishing it from other things that are to be said in some speaking of those words.

What those other features of a correct understanding might be is so far unknown. But we might conceive of them along the following rough lines: If the basic claim above is correct, then wherever something is said in saying, in terms of a more or less definite specification, by one means or another, of the circumstances or sorts of circumstances, in which the obtaining of what is described as being so in the antecedent, is said to go along (or concur) with the obtaining of that whose being so is described in the consequent. In different cases, different sorts of variations of the circumstances that actually obtained or the events that actually took place will be relevant for showing that the antecedent is a state of affairs that brings along with it the consequent, or that it is not. Which is to say, *inter alia*, that for the different things thus to be said in the given words, different sorts of circumstances and facts will be relevant to the *truth* of what was thereby said.

Suppose something like this is correct. Then, in special cases, it may be that the relevant specification will be such that what gets specified thereby turns out to be the circumstances that actually obtain, whatever they may be, and nothing else. In such special cases, there will, of course, be no suggestion that the consequent *follows* in any interesting way from the antecedent. Furthermore, what is thereby said (*ceteris paribus*) will be true if the antecedent is not, and true if the consequent is, and not otherwise. For example, suppose that the children are searching the house for Easter eggs, and Uncle Chester doesn't want them searching *too* hard

in his bedroom. Then he might say, 'Take my word for it. If there are any eggs in my bedroom, then they are under the bed', and no one would take him to have said or suggested thereby that there is something about eggs being in Chester's bedroom from which it follows that they *would* be under the bed. Nor would what he said be *false* if there turned out to be no eggs in the room at all. The existence of such cases ought not to surprise us. Nor does it provide any evidence at all against the thesis that what is *said* in speakings of 'if . . . then' is distinguished by extra (variable) features beyond anything which the meanings of those words in English might determine, and a corresponding great variety of distinct things to be *said* in so speaking.

We have not yet decided whether this thesis is true, or whether the appearances alluded to above are, indeed, nothing more than mere. Nor is it our purpose here to do so. What *is* of present interest, however, is what might establish, or at least support it. A good way to look for *indirect* support, at least, might be to ask this: What do we typically *do* with what is said in saying 'If A, then B? One thing, it appears, though not the only thing, is that we rely on, and use, what is thus said to be so in getting from premisses to conclusions, which would indicate that what is typically stated in such words is the facts, such as they are, that we do rely on in so moving. If we already know or accept as so that which is described in speaking 'A', and if we accept to be so what is said to be so in saying 'If A, then B', then we can, with the help of some general rules (*modus ponens*, for example), draw the conclusion that that which in the saying of 'B' was said to be so, is so (to take one example). In fact, the connection of 'if . . . then' to the drawing of such inferences seems at least as closely tied to the meanings of those words as anything else. But suppose this were not so. It would, anyway, be nice to be able to state those facts which, from time to time, we know to be so, and in virtue of which conclusions *may* be drawn from premisses. And it would be nice, in doing so, to state no more and no less than those facts of this nature that *do* obtain, and/or which we know to do so. If what there is to be said in the speaking of 'if . . . then' does not make that expression suitable for stating these facts, then we will need to find some other expression which *is* suitable for doing so. And, if none such existed, we would presumably need to invent some. That is, unless 'if . . . then' already does the job, we would have to introduce some expression, such as 'shmizex', which worked as follows: Where(ever) there is a fact or state of affairs consisting in its being correct to conclude, given that A obtains, that B obtains, that fact may be reported (i.e., *said* to obtain) in saying, 'A shmizex B'. The corresponding price, of course, is that such would be a mistaken report exactly where and whenever the reported fact did not obtain. One would not, in that case, have said to be so what is. The question is, how would expressions of the form 'A shmizex B' have to work in order to meet this condition? The path to an answer is that that depends on what facts or states of affairs there may be which consist in B's being correctly concludable from A. This depends, in turn, on how one such fact may be distinguishable from another (if any are

distinguishable from any), which depends, in turn, on where and when some such fact would obtain, and where it would fail to.

Suppose, then, that we *look* at the facts. What the few and sketchy examples above indicate is something like this: there is a fact of our getting to Breda being concludable from our being on the Bredaseweg which obtains where the Bredaseweg leads to Breda (and/or where we know, or are in a good position to think this), *and* such that it remains, and would remain a fact under circumstances where (a) there are many road names in the Netherlands consisting of a place name plus genitive plus 'weg' which name roads ending abruptly in cow pastures, (b) we are low on fuel, (c) our engine is making very funny noises—even where it is *not* concludable, given those noises, that the car will make it as far as Breda (whether the car will make it that far is not in question as far as the obtaining of this fact is concerned). There are even some such facts which would obtain even under circumstances where we in fact never get to Breda because, e.g., we lose interest, or the engine falls out, or we come to rest against a tree. Then there are other sorts of facts that would *not* obtain under some of these conditions—e.g., a fact that would not obtain given (a) above, but might despite (b) and (c), one that would not obtain given (b), one that would not obtain given that it is not concludable that our engine will last as long as Breda, one that would not obtain given that we have no way of working out that we are not going the wrong way on the Bredaseweg, and so on. A further note on *when* each such fact obtains: The first mentioned fact, for example, is such that on some (potential) occasions, or for some potential instances of (our) concluding B from A, it *is* correctly said to obtain, and for other such instances it is correctly said not to. And the same is true for all of the contrasting facts mentioned or alluded to above. As A-theses typically suggest, there is here again no question of making an inventory, *tout court*, of those facts of this kind which hold and those which do not.

Since, by hypothesis, 'A shmizex B' is usable for reporting *any* such fact—by saying it to be so—there will have to be many distinct things about A and B which are to be said in some speaking of 'A shmizex B'. Some report to be given in those words will be falsified, for example, where we are not in a position to conclude that the engine will last that long. Other reports—to be given in different speaking of the same words—will not. Whether a report *will* be thus falsified or not, it appears, will depend on how it is to be understood, in respects that are determined neither by the meaning of 'shmizex' nor by the proper understandings of 'A' and 'B'. It follows that 'shmizex' must function in such a way that a speaking of 'A shmizex B', where *some* complete thought is expressed, is to be understood in terms of extra distinguishing features—features over and above what the word 'shmizex' means, and over and above those which determine what was said in speaking 'A', and what was said in speaking 'B'—where those features contribute to the determining of what is to be understood as having been said in those words. Short of a specification of which of these features are

to be understood for a particular speaking, there will be many distinct things to be said in saying, 'A shmizex B'.

In short, 'shmizex' would have to function in exactly the way in which it *appears* that 'if . . . then' does function. This, I think, gives us good reason to believe (short of other candidates) that 'if . . . then' is our English version of 'shmizex'. Given that we are patently at no loss for words for stating the facts that 'shmizex' may state, and that there are no other candidates with notably better claims to the title of 'expression usable for doing so', we have once again found a reason to conclude that things *are*, linguistically, at least, as they seem.

On one reading of the facts, then, there is more to be said in a speaking of 'if . . . then' than can be determined by *any* correct account of what those words mean. Since there are many distinct things that this more may be, a truth-functional account of what is said, providing, as it does, only one such candidate cannot be correct. Grice, of course, reads the facts in a quite different way—one on which 'if . . . then' precisely *cannot* state the facts about what follows from what. But his reading requires one to take it that things are, in several respects, other than what they seem. On it, for one thing, where we appear to be speaking of what was *said* in some words, 'If A, then B', and correctly terming *it* either true or false, we *cannot* be doing that. Either we are *not* talking about what was said, or else if we are, our judgements about whether it is true or false are wrong. For another thing, on Grice's reading, contrary to appearances, 'if . . . then' is not really usable to say (as opposed to implicate) those facts to be so, in virtue of whose holding, the conclusions which are to be drawn from given premisses may be.

What sort of case *is* there, then, that appearances of the above sort *do* mislead? One thing Grice relies on is that if we propose a *particular* 'more that was said' in some speaking of 'if . . . then', then it is possible to construct some other speaking in which no such thing is even suggested—where, in his terms, that aspect of what was to be understood is *cancelled*. Another thing which he relies on is his ability to construct explanations of how the proposed 'more that is said' would come to be part of what was to be understood, even if its being so was not determined by anything about what the words 'if . . . then' mean. But we have already seen why neither of these considerations, even if true, supplies any evidence at all for Grice's position. The first does not, because cancellability is exactly what one would expect for any of that more which is to be understood as said, which may vary from speaking to speaking—that part of what is said, that is, whose identity is *not* determined by the meanings of the words used. And *that* something may so vary from speaking to speaking is no indication that it is not part of what is said. The second consideration provides no evidence because we know that there must, in general, be features and parts of what is said which are *not* determined by the meanings of the words used, and it is precisely such features that interest us as the 'more that is said' in saying 'if . . . then'. But for any such aspect of what is said, and for any given speaking for which that *is* an aspect

of what was said in it, there must be some facts about that speaking in virtue of which that aspect, rather than something else, is so. Insofar as Grice's explanations succeed, he may have put his finger on what those facts are. But doing so does not show that the aspects in question are *not* aspects of what was said.

How then *should* it be shown what the facts of the matter are? I suggest a sort of 'principle of transparence': Assume that people are saying or doing in their words (and especially in their words about others' words) more or less what they appear to be doing, and are prepared to recognize themselves as doing, unless there is some convincing reason to think otherwise. People *appear* to be saying different things in different speakings of given words, 'If A, then B' (with what is said in 'A' and 'B' fixed), and they *appear* (e.g., in their judgements as to what is true and what is false) to recognize themselves as doing so. Hence, there *are*, strictly speaking, different things thus to be said, unless there is good reason to think otherwise. So far, no good reason has been offered. In any case, the notion of conversational implicature by itself should provide cold comfort to anyone who wants to maintain that it is 'really' either true that A, or false that A, where an intelligent A-theorist has made out a case that we would not say the one thing, and we would not say the other. The onus is surely on the anti-A-theorist: What is it precisely that he feels must either be true, or not be? And, in precisely what words, and precisely when, is *that* what is, strictly speaking, said?

### On Working Things Out

Not that there is no distinction between saying something and implicating it, but Grice has not succeeded in making out what the distinction is—at least not in a way that helps settle any issues about what is and what is not actually said. Still, while some will be most interested in that distinction, and the uses for which it is meant, others may be more interested in the working out of what there is to be understood, whether said or implicated, and about what Grice has to say about how this might be done—in terms of what he terms 'conversational maxims', and given sorts of reasonings in terms of them. For, one may well have the feeling that these maxims are good for showing *something*, even if not for making out Grice's desired distinction. Let us now look at this feeling. First, *what* might the maxims show something about? Among the numerous possibilities are:

- a. What the speaker meant, what he meant or intended to say, what he was trying to say, or *was* saying, on the reading of this verbal aspect which entails neither completion nor success ('He was driving to Breda when his car hit a tree'), what the speaker was driving at, or insinuating, what point he was trying to make, and other matters only to be cleared up by and in establishing what the speaker's mental state was.
- b. What actually *was* said in a speaker's words—what, in virtue of a proper understanding of them, they, or the speaker in speaking them, actually said—a



matter which may run quite counter to what the speaker thinks, or meant, or meant to say, or in general to what is established in establishing his mental states.

- c. What there was to be understood as being communicated in given words or their speaking, either by being said, or otherwise—e.g., by being implicated—a question, like (b), not settled by what the speaker thinks, or tried to say or do.
- d. What in fact was communicated or gotten across, or what could be (or was) gathered from an utterance (or its uttering).

As subject matter for a theory of language, neither (a) nor (d) are particularly interesting, at least not in what they show about the properties of words, and what can, in virtue of these, be done in speaking them. Regardless, e.g., of what words mean, and of what might actually be said in them, a speaker may mean *anything* by them, or mean to say absolutely anything at all in using them, hoping to drive home, or insinuate, etc., any imaginable point. What he actually communicates depends on what his hearers may deduce, however irrationally, from his so speaking. From his words ‘Soybean futures traded lower in Chicago today’, one *might* deduce that he was in a sullen mood, or gather that his nth marriage was on the rocks, though he said merely that soybean futures were lower, etc., meaning to say in those words merely that it was raining in Singapore, thereby insinuating that Prince Charles’s interest in the Elgin marbles was merely feigned. Since pursuing such topics will not be very fruitful, let us concentrate on (b) and (c).

There are a number of sorts of theories one might seek about how what gets said gets said and ditto for what gets implicated. One might hope for conversational maxims to form the heart of any of these. We can, in particular, distinguish between predictive and retrodictive theories, e.g., of what is said in speakings. Given some words on a given meaning, either sort of theory might have something to say about the relation between a speaking of them and what was (in the retrodictive case) or what was or what would be, thereby, said. A predictive theory says, correctly if it is true, what will or would happen in every case (in this case, speakings) of which it treats. If its scope is speakings of English words as meaning what they do, then it will provide principles allowing us to derive from *some* facts about any such speaking what was, or would be said in it. It is concerned, that is, with what would be said in speaking words under the various circumstances in which they might be spoken (something which requires, of course, an adequate specification of what the *relevant* variations in such circumstances may be). It will say what is so, and nothing that is not—so, it will, for every case about which enough facts are fixed, propose a unique candidate—the right one. Further the input descriptions it requires must not already contain the results (or any part of them) that it is the aim of the theory to give. Input, that is, must not *entail* output, independent of the substantive predictive principles which the theory incorporates. Having such a theory would

tell us, thus, what it is in general about the uttering of words which makes it the case that what was said (or implicated) was so. A classic model of a predictive theory is, of course, classical mechanics. In terms of the initial motion of a body, and the forces acting on it, this theory aims to say what the resultant motion of the body will be, where nothing about this input information *entails* that that is what it will be, independent of the laws of nature which the theory aims to incorporate. That, after all, is how classical mechanics could turn out to be false.

A retrodictive theory is less ambitious. Either because of the complexity of the factors involved, or because some principles governing the interaction of these factors lie outside the scope of the theory, it does not pretend to *predict* what will happen in any given case within its scope—e.g., what *would* be said (implicated) in the speaking of given words, under given circumstances. However, given a particular speaking which *did* take place, and given what *was* said in it, the theory claims to provide principles, or a method by which, working backwards from this last fact, one can construct an explanation of how that came to be the case. Small-scale examples of nontrivial retrodictive theories are provided by medicine. If I am about to drink the water in a strange locale, a doctor may tell me, ‘I can’t tell you for certain that you will get sick if you drink this. But if you *do* get sick, I can tell you that it was the water.’ Whether I do get sick depends, perhaps, on the exact concentration of bacteria in the water, and the exact level of antibodies in me. Not only may the doctor not know this, but he may not even know any principles which would reliably tell him the outcome if he did. Nor, perhaps, *are* there any such principles. Still, *if* I get sick, it may be correct to say that we know that the level of antibodies was not high enough, that such-and-such bacteria caused my illness, and that that is what the doctor’s ‘theory’ correctly tells us.

Do conversational maxims have a future as the core of a predictive theory of what is said or implicated? It might be an interesting theory, if we would find one that was true. But our chances of doing so seem negligible. The *whole* story of why this is so is rather a long one. I have discussed pieces of it in other places.<sup>17</sup> Here, there is space for but an example of the problems which give cause for scepticism. First, a brief (and regrettably rough) synopsis of Grice’s maxims. Grice lists them under the headings ‘Quantity, Quality, Relation, and Manner’. In rough summary, they come, respectively, to: say what needs saying, and no more; don’t say what you don’t know, particularly not what you think false; be relevant (to the issues at hand, whatever they are); and, as much as possible, be brief, orderly, unambiguous, and avoid obscurity of expression. All of these injunctions have the character of rules we follow, and presume to be followed, except where we don’t. But the present object is not to question the validity of the maxims, but to see what they might explain. Let us, then, turn to examples.

<sup>17</sup> At least where this ‘term of art’ merely stands in for *ordinary* verbs of Grice’s intended class. For further discussion, cf. Travis (1981*b*).

Consider utterances of ‘A frog jumped in someone’s coffee yesterday’. Whether it is part of what is said or not, it will sometimes be to be understood, (i) that the someone in question is someone other than the speaker, and other times, (ii) that the someone in question might, or might not have been the speaker. For which speakings (i) is to be understood, and for which speakings (ii) is, is something for a predictive theory to tell us. For most speakings, at least, there will be no lack of things to say about maxims and their observance. For example, if the coffee jumped in was the speaker’s, then, often, there would have been a succincter way of being more informative, viz., by substituting ‘my’ for ‘someone’s’. In conjunction with the maxims of quantity and quality, this seems to provide good reason for taking (i) as that which is to be understood. On the other hand, (i) is *not* always what is to be understood. And there is no lack of facts about maxims corresponding to that fact, too. For example, consider the claim that frogs never jump in anyone’s coffee (though they may sometimes jump in unclaimed coffee). If that is an issue, then it is one to which what is said in the words as they stand is more relevant than is what would be said by substituting ‘my’ for ‘someone’s’. Perhaps that is sometimes a good reason to take (ii) rather than (i) as to be understood. But neither consideration can *always* be decisive in showing what is to be understood, since neither (i) nor (ii) always *is* to be. Sometimes (i) is, sometimes (ii) is, and sometimes it may be something else altogether. What our hypothetical predictive theory must tell us is for what speakings, or under what sorts of circumstances, the first sort of consideration has weight, and for what speakings the second does—and whether, and where this is weight enough to decide what was to be understood, and where not, what else *does* have such weight. In other words, the theory must identify those factors which may *vary* from speaking to speaking, and which make a given consideration sometimes relevant, and sometimes not. How might it do this?

The issue is not, I hasten to add, what considerations the *speaker* makes about the maxims, nor what *he* finds relevant or weighty, nor, in any sense, what he takes himself to be doing. What may thus go through his head is nothing for a theory of conversation, or of understanding to determine. Nor is it anything which *does* determine what such a theory must: the proper understanding of what, in fact, is said. Suppose, for example, that Max has been obsessed from youth with strange facts about animals, and in particular, with that old folk maxim ‘Frogs never jump in anyone’s coffee’ (a maxim known only to Max, as are his obsessions, though Max believes otherwise). Yesterday, a frog jumped in Max’s coffee, shattering for him a long-held and cherished belief. He has had little else on his mind since—except, of course, his present luncheon date with the recently met Clothilde. Trying, as he lunches, to say only what is most relevant to the topic uppermost in his mind, he utters our sample words, intending and thinking (ii) to be understood. Clothilde, who has an obsessive aversion to frogs, immediately changes the topic to fall fashions in Singapore, where it remains until she manages to be called away to the phone, never to return. Here,

though Max wanted to, he never got it to be to be understood that the person in question might just as well be him, which is to say that, not only was (i) in fact communicated, but it was also what was to be understood. One thing that would change this case would be if Max's obsessions were notorious enough. Clothilde, in that case, being sufficiently warned—as anyone who lunched with Max would be—it might well turn out that (ii) was to be understood.

But now, this compatibility of given facts about a speaking with various results about what was said or to be understood in or by it is exactly the problem for a predictive theory. For, if what, in fact, was in the speaker's head is one of the factors which the theory refers to in giving results in terms of the maxims, then the theory will also need to appeal to other factors in telling us when to take *this* fact as indicating one thing, and when to take it as indicating another (or nothing at all). And if those further factors are compatible with various end results—if, that is, they are just more of the same, as it is overwhelmingly likely that they will be—then the theory needs to refer to *yet* further factors to tell us how to take *those* factors, and when. And so on, *ad infinitum*. What seems to be the case, in short, is that *any* collection of independently ascertainable facts about a speaking (independently of how the words in it *are to be understood*, that is) is compatible with any number of distinct results about what there is to be understood, pending further considerations. In that case, the prospects for a true predictive theory, in terms of Gricean maxims, or in terms of anything else at all like them, appear exceedingly dim. This, of course, is not a proof, but a sketch of a reason for skepticism—the most that can be provided within present bounds.

What future do Grice's maxims have, then, as the core of a retrodictive theory? The question to ask here is not whether they might be incorporated in a *successful* retrodictive theory, since that would mean only that, by considering the maxims, we could, if we tried, always find *some* reason, in terms of their being followed (or not) for taking things as they are to be. The question is, rather, what such a retrodictive theory might possibly tell us.

Some retrodictive theories may well, if true, tell us quite a lot. One class of very self-consciously retrodictive theories which are at least interesting are psychoanalytic theories of such phenomena as the so-called transference neuroses, ordinary dreams, and various classes of common errors. (*One* interesting aspect of these theories, to my mind, is that they stand a very real chance of *not* being true.) One part of such theories is the postulation of certain specific happenings in early life. These happenings concern the maturation of a posited sexual instinct, and, concomitantly, specific early sexual activities, which are, in principle, largely open to direct observation. The maturation of the instinct is divided into stages, where each stage is characterized by specific observable activities, desires, beliefs, and kinds of imagining, of fantasies. The passage from each stage to the next is accomplished, in one or another or several manners, by specific psychological mechanisms. The last stage in this early period ends around age six (plus or minus a bit), and the whole period is succeeded by something else, called the

'latency period'. So far, the theory refers (or at least claims to) entirely to *historical* episodes and phenomena—concrete events which, if the theory is true, actually occur in early life, in the same way as the incurring of bumps, bruises, or chicken pox. These occurrences all may be, according to the theory, explanatory factors in the occurrence of certain phenomena in later life—dreams, for example—but only if they actually occurred, of course.

Psychoanalytic theories disavow predictive goals for essentially two reasons. First, there are too many intervening factors between age six, and, say, a dream dreamt at age thirty, for the theory to say either precisely what they are, or precisely how they may influence the effect of one of the postulated factors from early life. Second, the interaction of competing factors from early life takes place, in part, at least, according to principles which lie outside the scope of the theory—the dominance of one factor over another, for example, being governed, *inter alia*, by a factor which Freud calls 'psychic energy', the operation of which is not explainable within the terms of psychoanalytic theory. So, for example, a specific event that, on the theory, may take place in early life is that a child may experience something called 'oral frustration'. Nothing that the theory has to say entails that if this occurs, then certain things will happen in later life—e.g., the dreaming of such-and-such a dream, or the suffering of such-and-such a neurotic episode. If such things do occur in later life, however, then the theory has several things to say about them. First, if, e.g., oral frustration *is* the cause of phenomena in later life, then those phenomena are apt to have certain characteristics. So, from the features of a later phenomenon, one may be able to say, if the theory is correct, what the explanation of that phenomenon is likely to be. Second, the theory provides a method of working back from the phenomenon to its correct explanation (compare the case of working back from the illness to the drinking of the water). It also provides means for eliciting further phenomena for confirming or disconfirming a particular hypothesis. And third, in the case of an ongoing phenomenon such as a neurosis, once the explanation is arrived at, the theory (again, if correct) provides, in terms of it, means for altering the phenomenon.

A number of features of such a theory contribute to making it interesting. First and foremost, the explanations it gives, if correct, rule out a very wide class of rivals. One could, for example, posit a walking-upright instinct, or an eating-with-a-knife-and-fork instinct, with specific happenings in the development or frustration of these, and then try to provide, retrodictively, explanations in terms of these happenings for dreams, neuroses, and the rest. If psychoanalysis is correct, then these explanations are incorrect. And it cannot be known *a priori* which is which. Another such feature is that the explanations offered by psychoanalysis can, in principle, be tested, at least in the case of ongoing later phenomena, by using the explanations to manipulate the phenomena. Third, since the explanatory factors to which the theory refers are meant to be historical occurrences, they can, in principle, be observed to occur or not. And it is their actual occurrence or non-occurrence which, on the theory, matter.

All of the above appear to be missing in a retrodictive theory built around conversational maxims. Where such a theory tells stories, in specific cases, about reasoning that could be carried out, it does not claim that these are reasonings which anyone actually *does* carry out in coming to understand an utterance as he does. And for good reason: someone's actually going through certain reasonings or not on the occasion of a speaking will have no effect on how what was said was to be taken. Rather, it is the latter which determines whether his reasonings, whatever they may be, lead him right or astray. A consequence of this is that the success of a Gricean retrodictive theory would in no way rule out any rival theory as incorrect. And as a consequence of that, no such theory can make any pretension to offering 'the explanation' of what makes what is said (or implicated) in an utterance that which it is. Consider, for example, a rival retrodictive theory built around Travis's maxims of conversation (purely hypothetical, of course). Supermaxim: Always talk so as to leave the best impression of yourself possible. Submaxims: 1. Try to appear clever (or witty) where this would be appreciated. 2. Try, compatible with other maxims, to be maximally funny where you believe that your audience will prize humour. 3. Frame what you have to say in the most elegant and inventive way possible, where this is highly valued. (Subsubmaxim: be very brief where there is an aversion to verbosity.) 4. Speak earnestly and seriously where this is highly valued. 5. Speak modestly and self-effacingly, consistent with the other maxims, especially where you believe your audience to be averse to showing off. 6. Be as interesting as possible, and as inventive and imaginative as (but no more than) would be appreciated. (Subsubmaxim: Be bizarre, if you think your audience would go for that.) 7. Do not observe these submaxims where you wish not to make a good impression (e.g., to show your disdain), or where you believe that they would not achieve their intended effect.

In terms of some considerations or other about a speaker's presumed observance, or nonobservance of Travis's maxims, his presumed reasons for this, and what he would presumably take observance to be, we could doubtless always find *some* explanation of why one *might* take him to have said what he did. Why did Wu write that he had excellent handwriting when asked to write a reference for Mr Ng? Surely he knew that this wasn't directly responsive, and that thus giving nonanswers is apt to make a bad impression. There's no reason to think him to have acted in defiance of Travis's maxims, or at least the supermaxim. Perhaps he thought it would have made an even worse impression had he replied directly. If he simply said what he thought, he might have laid himself open for censure—perhaps as insensitive, or at least indelicate. And lying isn't really thought highly of either. Perhaps what he did made a less bad impression than that—in fact, by thinking of changing the subject in this way, when it turns out to be a painful one, he *might* even have thought to be thought clever. But why would the topic be painful? Most plausibly, he doesn't think highly of Mr Ng's abilities as a philosopher. So perhaps we can infer *that* from what he did. And if he knew this, which he must have, perhaps he was even suggesting that he

thought this about Mr Ng. One *might* in such lines of thought mention Mr Wu's presumed willingness, or lack of it, to be generally cooperative, etc. But one need not. And the chains of reasoning involved are neither the better nor the worse for the omission. Such success would in no way impugn the 'explanations' of the same phenomena which might be provided in terms of Grice's maxims. But nor would explanations in Gricean terms show Travis's maxims or the explanations they generate to be incorrect. Given such ecumenical tendencies, there is little reason to find either set of maxims interesting.

### Conclusion

It is easy enough, when confronted with Grice's maxims, to think thoughts of what they more or less fit. One asks whether Miss X is attractive, and gets the response, 'She's a very kind person'. Why would the speaker have said anything so unresponsive? If he didn't know, he might simply have said so: perhaps he did know, but would rather not say. And that could only mean . . . so we gather, in the imagined case, that Miss X is *not* so attractive. And it seems that we *might* do this in the way that Grice suggests we should. It is, of course, a triviality that *if* the respondent wished to be cooperative other things being equal—if, that is, he wished to answer our question—and *if* he understood that his answer did *not* answer it (he did not believe that kindness *per se* counts as what was meant by 'attractive'), then other things must not have been equal. And if the only thing that would make him give up his general resolve here was an unwillingness to admit that Miss X was not attractive, then he must have been unwilling to admit that—which means that he must have thought it. If that is the most probable explanation for what the respondent did, then it is most probable that that is what the respondent thought. If we have reason to believe all the above, then we have reason to believe that that is what he thought. So, from all of the above, we might gather that Miss X is not, in the respondent's eyes, attractive. We do not need Grice's principles, or any other special principles of conversation, to tell us this much.

Grice has provided us with maxims such that *if* we know that they are being followed, and *if* we know enough other facts as well, then certain things follow. That in itself hardly marks an advance. The question is how trivialities of the above sort are to be elevated to the status of a theory. What generalizations on such thoughts about particular cases would one wish to support? Which raises the question what such a theory is to be about. One possibility is a theory of what goes on in hearers' (or understanders') heads. I doubt that this is what Grice had in mind. But suppose someone does. Then what claims does the theory generate? There is no lack of claims on these lines which are clearly and easily testable, though none recommends itself as plausible. One claim is this: where a listener gets an answer which he regards as knowingly unresponsive—'She's

kind', 'She wears blue a lot', and/or so on—he assumes all of the above premisses and draws the above conclusion. Such must be tested against cases like the following. Suppose the hearer knows that the respondent is angry with him, or being playful, or a relentless supporter of the theory that beauty is only skin deep, and not to be asked after. Will he still, despite this knowledge, draw those conclusions? Perhaps so, but it hardly seems likely. But perhaps the theory is not intended to predict this. Perhaps it is only designed to tell us what hearers will do when they take it that speakers are following policies entailing, *ceteris paribus*, such things as being responsive—Grice's maxims, for example. Even so, it is hardly plausible that when people are unresponsive, hearers suppose them to be so because of reluctance to admit the truth, as opposed, say, to disapproval of the interests manifested in the question, or desire to take a shortcut to what they suppose really to be on the speaker's mind. Nor would one want a theory of conversation to entail any such thing. We might then try this: Perhaps what we want our theory to entail (and no more) is that where a hearer takes a respondent to be knowingly unresponsive, he will gather from that whatever he takes to be the best reason for that unresponsiveness, given that there is at least some plausible enough candidate. If the best reason he can think of is that the speaker is reluctant to admit that Miss X is not attractive, then he will conclude that Miss X is not, in the speaker's eyes, attractive. And if there is reason to think in the conversation at hand that the speaker *would* be adhering, *ceteris paribus*, to policies of responsiveness, etc., then this *may* well (though it need not) substantially increase the chances that this last *will* be the best available explanation for what the speaker did. And if we build in enough conditions about what the hearer must suppose before the theory is to apply, then we will quickly arrive back at trivialities, rather than hearer psychology, or any substantive generalizations on principles of conversation. In any case, theories on these last lines may stand a large chance of being true, given that people generally draw conclusions in line with what they take to be the best reasons for happenings (if they do). What such a theory doesn't tell us is what someone *will*, in fact, conclude in any particular case, much less what was implicated—where in such responses, the above *would* be implicated, and where not; where something else would be (e.g., that Miss X *is* attractive, kindness being the proof thereof, or that the speaker was impatient with our prolixity—taking it that kindness was what we, in a round about way, were trying to get at), and so on.

But, given the discrepancies between what hearers *will* do, and what *was* implicated (given that 'implicate' is an umbrella word for 'imply', 'hint', 'insinuate', and an unspecified variety of others), perhaps we may conclude that the aim is a theory of another kind. Perhaps the claims of the theory concern what is to be (as opposed to what in fact is) gathered—e.g., from presumptive unresponsiveness. But a theory along these lines will run into roughly similar problems, among others. Ought one really to gather that the speaker finds Miss X unattractive when it is well known that the speaker believes that it takes a very attractive



person to be kind (or rich, or whatever), and there is every reason to take him to be capitalizing on this fact? Or when it is clear enough that he simply doesn't want to let *us* know what he thinks of Miss X? One can build into the claims of such a theory that they are only to hold when it is to be gathered that such counter possibilities are not the case—e.g., when it is to be known that the speaker does not have other special reasons for desiring to be unresponsive—but that we may only draw the above conclusion when it is to be taken that *that*, rather than something else, is the reason for the speaker's unresponsiveness. But then we are back to trivialities which we need no theory of conversation to tell us.

Perhaps, finally, one really wants a theory of a more retrodictive character: one that says what goes where one *does* gather, or where it is to be gathered that in the speaker's eyes, Miss X is not attractive—where, e.g., anyone would gather that with little hesitation. Again, there are various possibilities for what one says *does* go on. Is it to be reasonings in fact carried out in the hearer's head? Or is it to be that in such cases there will have been certain things to be taken to be so—e.g., that the speaker had such and such a reason for nonresponsiveness? But again, any such theory, no matter how framed, is on pretty thin ice. Must there, for example, have been reason to believe in the speaker's reluctance to admit the truth about Miss X where, in fact, it is known that he is a great admirer of ugly women, or a great enemy of Miss X? Or where *he* is known to be more concerned with Travis's principles than with Grice's? Once again, when we look beyond the cases that occur to us at first hearing of Grice's maxims, it is hard to see how to make such a theory plausible without reducing it to triviality.

As the nucleus of a theory of why conversation is what it is, conversational maxims suggest nothing plausible when viewed so as to be interesting, and nothing interesting when viewed so as to be true. Perhaps it is viewing them in both ways at once that creates a spurious image of some philosophic or linguistic purpose being served in speaking of them.

## 2

# Annals of Analysis

H. P. Grice transformed, often deeply, the problems he touched and, thereby, the terms of philosophical discussion. The discussions he touched most deeply concern meaning, saying, and perceiving. All Grice's work displays a particular method, practised with consummate skill. Often, I will suggest, the method decides much as to what his view must be. *Studies in the Way of Words*, a self-selection of his work, provides an occasion for evaluating both the transformations and the method they represent.

### 1. Implicature

There is an intuitive distinction between what was actually said in given words and what was only suggested, implied, meant or being driven at. For example (due to Sperber and Wilson 1986), if, asked whether I want coffee, I respond, 'Coffee would keep me awake', then I have not *said* that I do not want any. Perhaps, though, what I did say is properly to be taken to convey that. For some of what is not said, but otherwise conveyed, that it was to be gathered is part of the meanings of the words used. '*N* has stopped beating his wife', on any speaking, would imply that *N* once beat his wife. In such cases, we might speak of presupposition. (Grice prefers 'conventional implicature'.) In other cases, what is not said, but is otherwise conveyed, depends, for its being conveyed at all, on occasion-variable factors: it is implied, suggested, or etc., by/in those words as spoken in those circumstances, but would not be suggested by the same words, spoken as meaning the same, but in other circumstances. Grice refers to what is otherwise conveyed in such an occasion-variable way as *conversational implicature*.

There are, no doubt, a pair of distinctions of the sort just sketched. One might question Grice's uses for them: to deflate philosophical theses of a certain genre. To think them so usable, I will suggest, is to mistake what those theses are. Here is a sample of what Grice targets:

Malcolm accused Moore of having misused the word 'Know' when he said that he knew that this was one human hand and that this was another human hand; Malcolm claimed, I think, that an essential part of the concept 'know' is the implication that an enquiry is under way. Wittgenstein made a similar protest against the philosopher's application of the word 'know' to supposedly paradigmatic situations. (p. 5. For Wittgenstein's protest, see *On Certainty*, §6)

Other theses targeted by Grice involve trying (Wittgenstein) ('You can't try if there is no difficulty'); Austin's maxim 'No modification without aberration', involving 'voluntarily', 'deliberately', etc.; 'seeing as' (Wittgenstein) (one does not invariably and ineluctably see a tree as a tree); theses about connectives and quantifiers (Strawson: 'or' is not truth-functional); and theses about illocutionary functions of certain words (e.g., 'it is true' expresses agreement or concession, 'I believe' expresses reservation or doubt).

To exhibit Grice's way with such theses I concentrate on knowledge. Here Grice has his opponent maintaining some such thesis as 'One would not say "I know that *P*" unless there were some doubt as to *P*', which illustrates the general form of thesis Grice ascribes to his opponent: 'One would not say *A* unless *B*'. The general response to such a thesis is: perhaps we would not (usually) say *A* if not *B*; but perhaps for the following reason: to say *A* where not *B* would be (conversationally) to *implicate* something that was not so—in many cases, that *B* itself (for example, for knowledge, that there is some doubt). For all that, if one *did* say *A*, what one would thus have *said*, as opposed to implicated, might be *true*. So a philosophical theory which represents such things as true (e.g., Moore on knowing one has hands) may survive the admission: Moore may have spoken truth in saying he knew he had a hand, even if we would not normally say such things (where there was no doubt).

Grice does show that the bare remark 'We wouldn't say *A* unless *B*' often leaves philosophic issues untouched. His strategy misfires, though, against the theses of a number of actual philosophers, notably Wittgenstein and Austin, which are *not* of that form. It founders there on Grice's simple view of saying—one already present in his tests for the presence of implicature. Of these, the most impressive, cancellability, amounts to this. Suppose that words *W*, spoken on an occasion (or their speaking then) conveyed that *P*, whether by saying it or by other means. Suppose that the same words, spoken literally on other occasions, would convey no such thing. (For example, reviewing suspects, 'We know that the butler had the day off' suggests no doubt about it.) Then, Grice thinks, this is evidence that that *P* was not said, but otherwise conveyed. Now, that would be strong evidence *if* we supposed a certain link between what words *mean* and what is *said* in them on an occasion—roughly, one on which words which meant what *W* do could only *ever* (literally) say that *P* if they always do so. Otherwise, it is no evidence at all.

Though the issues here are largely beside the point for the actual theses Grice sets himself against, it is worth following up this clue to his view of saying. Grice discusses what is said, by way of introduction to implicature, as follows:

In the sense in which I am using the word *say*, I intend what someone has said to be closely related to the conventional meaning of the words (the sentence) he has uttered. Suppose someone to have uttered the sentence *He is in the grip of a vice*. Given a knowledge of the English language, but no knowledge of the circumstances of the utterance, one would know something about what the speaker had said, on the assumption that he was speaking standard English, and speaking literally. One would know that he had said, about some particular male person or animal *X*, . . . . But for a full identification of what the speaker had said, one would need to know (a) the identity of *X*, (b) the time of utterance, and (c) the meaning, on that particular occasion of utterance, of the phrase *in the grip of a vice*. (p. 25)<sup>1</sup>

For Grice, then, what words mean very nearly determines what would be said on any speaking of them (as meaning what they do); so what their truth or falsity so spoken would consist in. In fact, what words mean fully determines one thing which is what they say, and would say on any speaking, modulo a few specifiable speaking-variable factors, where what these are, and *how* they vary—how fixed for each speaking—follows from what those words mean. If, e.g., the meanings of words show them to be speaking of a certain time—say, the time or place of their speaking—then that time is part of (a specification of) what was said in an instance of them, and must be supplied by their speaking. That it must be follows from what those words mean. Similarly, a speaking must supply contextually indicated objects for words whose meanings show them to speak of ones on a speaking, and it must remove lexical or structural ambiguities, where meaning shows them to exist, as in ‘Mary had a little lamb’. These indicated factors are simple, few, and transparent to one who understands the words which call for them.

Moreover, ignoring linguistic ambiguity, what the words of a sentence mean always allows us to specify, by stating it, which thought a sentence would express, for any speaking of it. Suppose the sentence is ‘The table is covered with butter’ (Grice’s example). Suppose that what its words mean shows it to refer, on a speaking, to an indicated object (a table), and a time (of the speaking), and to demand no more factors to be supplied by the circumstances. Then the thought expressed in that sentence, on a given speaking, *S*, is the thought that the object (table) indicated in *S* was, at the time of *S*, covered with butter. This illustrates how, in the absence of any details about a speaking, ‘that’- phrases are apt: for any non-ambiguous sentence, *S*, we can always identify at least one true filling of

<sup>1</sup> Relying on British spelling, Grice finds this phrase ambiguous between ‘caught in a certain sort of tool’ and ‘a slave to a bad habit’.

the blank in 'It was then said in *S* that \_\_\_\_\_'. We may also correctly state when what was thus said would be true—e.g., iff the indicated table was then covered with butter. Such richness of semantic fact is derivable simply from what *S*, or its parts, *mean* (in their language).

Grice's view of saying has a corollary for the theses he aims to undermine. On Grice's understanding of them, they take the form 'We would not say *A* unless *B*'. Perhaps not. But, on Grice's view, if we did say *A* where not *B*, we would clearly have expressed a thought, and it would be clear which thought it was; thus what it would be for truth to have been spoken. The only proviso is that it remain clear what the time and place of the speaking were, what the relevant indicated objects were, and so on. Given this, it is difficult to see what such a thesis could be saying, other than that saying (that) *A* where not *B* would be odd, eccentric, misleading or bad manners. It is hard to see how the thesis could concern the sort of thought, if any, expressed in *A*, or what such thoughts require for truth.

Grice's simple view of saying contrasts strikingly with that of his opponents. Austin (1962*b*), for example, says:

. . . if you just take a bunch of sentences . . . impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for . . . the question of truth and falsehood does not turn only on what a sentence *is*, nor yet on what it *means*, but on, speaking very broadly, the circumstances in which it is uttered. (p. 111)

A sentence (nearly any) may, on one speaking or another, say any of indefinitely many distinct things, each true under different conditions. Nearly any part—a simple predicate like 'is red', say—may make any of many contributions to what it thus says, specifically to the conditions for its truth. *All* these contributions are ones those words would sometimes make *given* what they mean; all are compatible with their meaning that. Nor does their meaning provide the means for deriving when they would make which contribution. Rather, seeing what words did, or would, say on a given occasion is a matter of properly appreciating the circumstances of that speaking, and correctly perceiving which of their many possible contributions they are most reasonably taken to have made in those circumstances.

A token of this changed perspective is that homophonically constructed 'that' locutions are no longer *automatically* apt for specifying what given words said. Suppose the words were 'The table is covered with butter'. May we say that they said that the indicated table was covered with butter? Only if *our* words 'was covered with butter' make the same contribution to their whole as the described words 'was covered with butter' did. But consider. Suppose I buy masses of foil-wrapped packets of butter. I then arrange them on the table so that no bit of surface is showing. Is the table covered with butter? Intuitively, there is something true to be said in so describing it, and also something false. What one

would say in such words, equally, what the words quite literally *said*, whether a true thing or a false one, depends on such things as why one is saying it. On Austin's view, such intuitions are often strictly speaking correct. If so, then 'was covered with butter' makes different contributions to its whole, and what is said in it, on different occurrences, or speakings, of it. If reported words made one sort, while those in our 'said that' report made the other, we did not speak truth. Nor could we guarantee success by getting *our* 'was covered with butter' to make neither contribution.

By the same token, what words mean gives us no automatic access, as it would on Grice's view, to what is required for their truth as spoken. Is Pia's 'The table was covered with butter' true iff that table was then covered with butter? Any of many conditions for truth might be stated in so speaking, depending on which of its many possible contributions *our* 'was covered with butter' made to the truth condition thus stated. Which of these conditions is correct depends on more than we so far know of what was said in Pia's words. Nor is what they mean enough to make us any wiser on that score.

Wittgenstein (1953) presents the same picture:

You say to me: 'You understand this expression, don't you? Well then—I am using it in the sense you are familiar with'.—As if the sense were an atmosphere accompanying the word, which it carried with it into every kind of application.

If, for example, someone says that the sentence 'This is here' (saying which he points to an object in front of him) makes sense to him, then he should ask himself in what special circumstances this sentence is actually used. There it does make sense. (§117)

The issue is one of *making sense*; not one of what we wouldn't say. (Current accounts of indexicals obscure the point.) The point is repeated in §501:

'The purpose of language is to express thoughts.'—So presumably the purpose of every sentence is to express a thought. Then what thought is expressed, for example, by the sentence 'It's raining'?

And again in §514 and §515:

A philosopher says that he understands the sentence 'I am here', that he means something by it, thinks something—even when he doesn't think at all how, on what occasions, this sentence is used. And if I say 'A rose is red in the dark too' you positively see this red in the dark before you.

Two pictures of a rose in the dark. One is quite black; for the rose is invisible. In the other, it is painted in full detail and surrounded by black. Is one of them right, the other wrong? Don't we talk of a white rose in the dark and of a red rose in the dark? And don't we say for all that that they can't be distinguished in the dark?

Is 'The rose is red' true of a rose in the dark? On one picture of a rose's being red, yes; on another, no. Depending on their surroundings, those words, by what they literally say, might present us with either picture as that in terms of which they are to be evaluated. That is why we should not think of them in isolation,

nor of 'I am here' in isolation—even with an indicated speaker and time—as expressing a thought; saying something to be so which is determinately so or not. There are many mutually conflicting, sometimes correct, ways of counting things as being red. There are, correspondingly, many things, each sometimes to be said of a rose, in saying, 'It's red', where 'is red' speaks of being red, and means what it does mean.

As shown by the contrast with Grice, Austin and Wittgenstein introduced an essentially new view of the relation of language to what is said in using it (to thought). It changes the form of most philosophic problems. The new view concerns the relation between words and the various semantics they may bear—what they may say, conditions for their truth, and so on. On the old view, with a few explainable exceptions, a word or expression of English does what it does on each use of it in virtue of a semantics it bears as such, hence on each use. This semantics is rich enough to confer on words such roles *in their language* as expressing such and such thought, or being true of such and such. Standardly, corresponding to an English sentence, there is the thought it expresses; a sentence typically says *that* thus and so. (A sentence with indexicals could be thought of as expressing a thought-schema, convertible to a thought by adding to the schema the features of the speaking it indicates: its time, place, etc. The 'says that' locution would still fit: 'The table is wet' would say that some indicated table is, at the time of its speaking, wet.) This view allowed Frege (1918–19: 60) equate thoughts with the senses of some declarative sentences. More recently, it has allowed the illusion that the meaning of an English predicate fixes a function from objects to truth values.

The new view might be called the speaking-sensitive view of words, or language. On it, words are sensitive to their speakings in the semantics they bear, varying semantics across speakings. So any semantics they might bear in saying something to be so is one they bear only occasion-sensitively. Their semantics as part of their language, e.g., English, is at most a proper part of their semantics on an occasion of expressing a thought, and underdetermines what thought they would thus express, the latter varying while they mean what they do and have whatever semantics that confers on them. Their fixed, language-contributed semantics must, in general, be supplemented if they are to be properly assessable as to truth; that is, if they are to count either as true or as false. Nor is such supplementation provided automatically by some fixed stock of 'indexical features'.

The role of a sentence, on this view, is *not* to be the expresser, in its language, of such-and-such thought, but rather to be usable in many different circumstances for expressing any of many thoughts, each with its own condition for truth. The view thus respects a gap between what words *mean* and what is said in them on occasion. We, or our words, may say *that* thus and so. Those words do not *mean* that. Nor do they mean *that* that, or that anything else. Nor, where we, or they, say that thus and so, is what is thus said derivable simply from what they mean.

With this difference between Grice and his opponents in mind, let us return to the specimen debate. On a Gricean view of saying, Moore's 'I know I have hands', while perhaps bizarre or inept in any of many ways, must nevertheless have said either what is true or what is false. Even though, perhaps, we would not say such things, if we did say them, as Moore did, we would express truth, or, *casu quo*, falsity. Thus, Malcolm's notion of 'would say' is a vicious fudge factor. (Contrast this view with Wittgenstein's view of a rose in the dark.)

A speaking-sensitive view would oppose this Gricean view in two stages. The first would consist in a demonstration that the word 'know', used strictly and literally, may make any of an indefinite variety of distinct contributions to what is said in speaking it, and, specifically, to the truth condition for that. (Precisely what these contributions are will depend on the precise nature of the occasion-sensitive account on offer.) Correspondingly, on the account, many distinct, independent, states of affairs may each sometimes count as *N*'s knowing that *F*. Each of these is what would be said to obtain in *some* words, '*N* knows that *F*'. What any one of these says to obtain differs from what any other says to obtain.

The second stage consists in drawing a corollary from the above point. Words '*N* knows that *F*', to be either true or false, must speak of some one or another of those various states of affairs each of which sometimes counts as *N*'s knowing that *F*, and say exactly that to obtain. For them to do that, their semantics must be a supplementation of the semantics they carry to each occasion as words of English. What they mean in English does not choose any of these states of affairs in particular as the one they say to obtain. Moore cannot rely on the mere fact that we all know what his plain English words mean, or that they mean what they do, for showing, or for fixing, what he said (to be so) in his words. He cannot rely on meanings alone to ensure that he expressed a thought at all—that is, said what some states of affairs, and not others, would make true. Whether anything is said in words like Moore's, and, if so, what, must be fixed by the surroundings, or facts, of their speaking; his words will bear that supplementation of their constant semantics, if any, that those surroundings make most reasonable. But surroundings of some speakings may be inadequate to this semantics-fixing task: they may fail to choose between competing things to be said in those words, where some of these competitors conflict in their conditions for truth, some conditions being satisfied, others not. In that case, what those words then said, if anything, could not be something true, nor could it be something false.

Moore had no particular purpose in saying 'I know I have hands', except to produce an example of something he knew. The result of that, one might argue, is that the circumstances of his speaking were inadequate in the way just described: they did not choose a supplement for the semantics of his *English* words rich enough to fix anything possibly either true or false as that which was said in his words. The result would be that Moore, speaking as he did, said nothing



true, and nothing false. He said nothing to be so; so nothing one might refrain from saying. Moore did *not* state something most of us would keep quiet about, perhaps because it is bizarre or impolite to mention it. This contrasts sharply with a Gricean view. It shows that if we take speaking-sensitive views seriously, Grice has not yet demonstrated that it is beside the point if we would not say *A* unless *B*.

To say 'We wouldn't say *A* unless *B*' *may* be to say: circumstances in which not-*B* are inadequate for fixing anything sufficiently determinate as that which was then said in saying *A*, given the variety of distinct things that may be said in saying *A*, speaking literally. (Such may be enlightening, even if no more than a rule-of-thumb.)

If *B* thus functions as a condition on a speaking's surroundings, required for them to fix something sufficiently determinate as that which was said in saying *A*, there is no presumption that *B* is part of what is *said* in saying *A*. Saying 'I know' is not *saying* that there is some doubt about it. Nor, where something is said to be so in saying *A*, need *B* be part of what is required for the truth of that. Rather, *B* is simply a condition on saying *anything* to be so in saying *A*; part of what is required for surroundings to do the work required of them.

Aiming at accounts like this, Grice spends most of his ammunition trying to show that such-and-such would not be *said* in given words—e.g., to say that *N* knows that *F* is not to *say* that there is doubt as to *F*. That is why he places special emphasis on remarks like the following:

If I say (smugly) *He is an Englishman; he is, therefore, brave*, I have certainly committed myself, by virtue of the meaning of my words, to its being the case that his being brave is a consequence of (follows from) his being an Englishman. But while I have said that he is an Englishman, and said that he is brave, I do not want to say that I have *said* (in the favoured sense) that it follows from his being an Englishman that he is brave, though I have certainly indicated, and so implicated, that this is so. I do not want to say that my utterance of this sentence would be, *strictly speaking*, false should the consequence in question fail to hold. (pp. 25–6)

Showing such things is the main use to which the notion of conversational implicature is put. In so deploying it, Grice merely jousts at phantasms. The philosophic views that implicature is aimed at do not require things to be said *in addition* to what Grice would want to recognize. Rather, they challenge the idea that, in those sensitive cases, Grice has identified anything to be said at all, or anything of the sort we ever do say in speaking of relevant topics (e.g., knowledge). They may posit needed background for expressing relevant thoughts: conditions we do not speak of in saying things on the topic, but which must obtain for us to have said anything on the topic at all.

Suppose that we normally would not say '*N* knows that *F*' (e.g., 'Sam knows she is wearing shoes') unless *B* (e.g., there is some specific sort of ignorance one might be suspected of). Grice's explanation of that fact would be either that,

normally in saying what one would thus say, one would implicate  $B$ , or that in saying that where not- $B$ , one would implicate  $C$ , where  $C$  is obviously false, or anyway not what one meant to communicate. As it happens, Grice insists that it is a fact that  $N$  implicated that  $B$  *only* where that fact is derivable deductively, by appeal to his conversational maxims, from the fact that  $N$  *said*, strictly, literally, that such-and-such else. So that form of explanation is available to Grice only where he can specify a suitable candidate for that which  $N$  *said* in those words, and what it would be for that to be true. Part of the force of Wittgenstein's and Austin's epistemology is that no viable such candidate is available; the only non-arbitrary candidates being things we would never say, or want to, or take ourselves to—things that actually *would* be refuted by the mere existence of sceptical doubts. Whereas the fact that the sceptic can produce *some* such doubt comes as no surprise to anyone. To see us as saying such things would exhibit poor method in linguistics; a bad fit of theory to data.

Grice's account of how implicature arises could be changed. Perhaps one can implicate things merely by *trying* to say something, or merely by speaking given words, whether or not one thereby says anything, or by remaining silent, or making a face, or by anything. That changed view of implicating would change the consequences of  $X$ 's having implicated  $B$  in saying ' $N$  knows that  $F$ '. It would not follow that  $X$  would have said that such-and-such anyway, whether or not  $B$ , or that, if  $X$  did thus say something to be so, the truth conditions for what he said are independent of  $B$ . It might be, for example, that the condition for the truth of what  $X$  said makes essential reference to the specific doubts which were to be taken to be real for  $N$  as to  $F$ .

A speaking-sensitive account of knowledge needs to be established. But since it is at least a viable form of account, adjudicating between it and a Gricean one may proceed by normal standards of good linguistic theory. The occasion-sensitive theorist would assemble ranges of cases, in all of which someone said of  $N$  (Pia, say) that she knew (didn't know) such-and-such (that Sam plays chess, say), and in each of which what was thus said is intuitively different from what was said in any of the others. The cases might differ, for example, in when we would intuitively regard what was said as true. The speaking-sensitive account would take these intuitions at face value. The Gricean account, for its part, would produce *one* candidate for what was said in all these cases. So it would represent us as often *literally* saying things quite different from what we ordinarily take ourselves to have said. The implausibility of that construction of the data would be softened to some extent by the observation that in such cases we often implicate what we intuitively treat ourselves, and each other, as having said. The Gricean might also build a case that his account is truer to what we are all prepared to recognize about what words like 'know', as such, *mean*. The question would then just be: which account handles all relevant data most plausibly? *Is* it plausible, for example, that our intuitions as to what we say are discountable, in the Gricean way, as merely about what we

implicate? Or that quite mundane words are typically used by us to state literal falsehood?

The Gricean faces two major hurdles. First, given a wide range of candidates for things sometimes to be said in '*N* knows that *F*', the Gricean must argue that his candidate, selected from, or added to, the range, genuinely has a better, more reasonable, claim to be what is always literally said in those words than anything else in the range; that his choice is not arbitrary. Many candidates obviously fail that test. Suppose, for example, that two candidates for what is said in 'Pia knows that Sam plays chess' differ only in that the one, and not the other, is false if Pia has not verified that the person she saw playing chess is not a ringer, made up to look just like Sam. Obviously, neither is a better candidate than the other for what is *always* said in those words. So there is considerable pressure to choose for the Gricean role either something insensitive to all doubts as to *F* (so that Pia knows or not, regardless of what she has or hasn't checked), or one that treats all doubt as real. But both of these sorts of candidates are monumentally implausible; the latter precisely because it yields scepticism.<sup>2</sup>

The other hurdle is that in all *clear* cases of implicature—all the cases Grice and others use to get that notion off the ground—what is implicated in given words is something totally different from, and an addition to, what is said in them. Pia says that she has just eaten, and implicates that she would not like to lunch with Sam just now. There is no clear philosophically innocuous case where what is said is that *X* is *Y* in *one* sense of being *Y*, while what is implicated is that *X* is *Y* in another sense of being *Y*, or where such-and-such counts as being *Y*. In any such case, the suspicion would be strong that there are just different things to be said in saying something to be *Y*, different things all sometimes (rightly) called 'being *Y*', thus no real implicature at all. In the philosophically interesting cases, where implicature is a tool employed against some speaking-sensitive account, that is almost always what the saying/implicating contrast comes to. Such use of the notion gains no plausibility from its clear instances.

A speaking-sensitive account of knowledge explains how philosophical perplexities arise: they do so naturally when 'language is on holiday'; when we are not speaking in surroundings where we would actually express a thought in saying *N* to know *F*, but suppose that we must be expressing one anyway. It thus also details precisely what misunderstanding scepticism is. Similar remarks apply to other philosophical appeals to speaking-sensitivity. The point is to give a new reading to intuition, thus reconstruing the impossible, conflicting requirements it seems to impose on some philosophically crucial concept. Though conversational implicature is a valuable linguistic notion, Grice deploys it so as to mask genuine insights which speaking-sensitive approaches supply.

<sup>2</sup> Some see hope that nature will sort out the real from the mere in a speaking-insensitive way—e.g., by making some other 'possible worlds' the ones we are actually close to. Within present confines, all I can say to such naturalism is: Good luck!

## 2. Causes and Perception

In ‘The Causal Theory of Perception’ (essay 15), Grice defends, with qualifications, an account on which a necessary condition for seeing  $X$  is that there be a causal link between  $X$  and something. His aim is to specify the required something. Grice’s suggestion is anachronistic in ways he was aware of. Much of the work of the essay is an attempt to defuse the arguments that make it seem dated. As Grice states, those efforts are the original impetus to his views on conversational implicature. I will not try to assess the correctness of Grice’s proposal. I aim only to make clear the sort of proposal it is, and how Grice’s conception of philosophy pushes him towards it.

Seeing requires, on Grice’s view, a certain sort of causal link. On one end of the link is the thing seen—an apple, say. Grice’s proposed occupant of the other end of the link is: a sense datum. One must add immediately: the notion of sense datum in play is to be a sanitized one. Grice recognizes that *sense datum* is a notion in some disrepute. His project is, first, to show how to introduce the notion so as to avoid the justified sources of this disrepute; and second, to defuse what he considers to be unjustified charges against it.

To his credit, Grice does not merely introduce a new term for sense datum, say ‘percept’, and *assume* he is speaking of something when he uses it. He undertakes to show how his terminology is to be used, and that, so used, it is coherent. But before approaching those details, let us observe what Grice wants sense data for. It is a natural thought that if I see something—an apple, say—there is some sort of causal interaction between me, or my visual apparatus, and the apple; and that if there were not, it would not be *seeing*. What sort of interaction might be required? There are the sorts of causal relations between the apple and my eye or brain about which we may hope to learn from physics and physiology. These do not interest Grice much.<sup>3</sup> Nor would his story change essentially with them in the picture. He would still need sense data in addition, for the reason about to be stated.

What might the apple cause? Perhaps: me to see it. Or its being there might cause that. Such suggestions, though, do not serve Grice’s purpose since he seeks an *analysis* of seeing. His aim is to state what it is in general, for any  $X$  and  $Y$ , for  $X$  to see  $Y$ :  $X$  sees  $Y$  iff . . . . But ‘ $X$  sees  $Y$  iff  $Y$  causes  $X$  to see it and . . .’

<sup>3</sup> It is worth noting that Grice has very bad reasons for his lack of interest in this sort of causal chain. He says, ‘if we are attempting to characterize the ordinary notion of perceiving, we should not explicitly introduce material of which someone who is perfectly capable of employing the ordinary notion might be ignorant’ (pp. 239–40). The person in the street may be ignorant of physiology, as we all are in present matters, the right theory being still under development. Still, it might be part of the ordinary notion that there is (to be supposed to be) some crucial physiological difference, to be discovered later, between seeing and not; that the right things, whatever they may be, must occur physiologically if it is to be seeing, where there are particular such things that do go on and do make something a case of seeing. On this point we have learned a great deal from Hilary Putnam (the *locus classicus* being Putnam 1975).

cannot constitute an *analysis*—unless a circular one, understandable only if we grasp what *Y*'s causal effect is to be, so, what seeing is.

Grice's goal here, I think, is part of his conception of philosophy: in large part, finding or testing analyses of (certain) concepts. Suppose we abandoned analysis. (We might think seeing Butlerian, that is, not anything *else*, hence admitting no analysis in Grice's sense. Recent studies, notably by Hilary Putnam, of the factors that fix what words would apply to suggest as much.) Having abandoned *analysis* of seeing, the remark that *X* saw *Y* only if *Y* caused *X* to see it, if true, might be of use. Perhaps it can be shown in particular cases, without having decided whether *X* saw *Y*, that *Y* could not have caused such a thing; so, that *X* could not have seen *Y*.

Seeing an apple is not plausibly *mere* physiology. It is, in any given case, a particular experience. In an *analysis* of seeing, this experiential component must be given its due, at least in Grice's view. However it so figures, though, the relevant experience had better not be seeing an apple. So we need some *other* experience such that having that under the right conditions *is* seeing an apple.

It is now clear what form Grice's account *must* take. Modulo the role of physiological termini, the form is forced by the fact that analysis is the aim. It is this: for *X* to see *Y* is for *X* to have a particular (visual) experience, *E*, and for *E*, or *X*'s having *E*, to be caused by *Y*. Some questions remain open. There might be a particular experience, *E*, one always has in seeing an apple; or *E*, in each case, might only need to be of an appropriate sort. (It could not be like seeing angels under a rosebush.) What is not an open question is this: *E* must be distinct from, and independent of, my seeing the apple. If I had *E*, and *E* had the right causal antecedents, then I saw the apple. But one might have *E* without seeing the apple (e.g., in hallucinating). If I had *E* in the wrong way, I might not, while having *E*, be aware, or able to be, that that was so; it might seem just like, or as if I were, seeing an apple. *E* is a fit component for deceptive cases.

The role sense data are to play in Grice's account is also now fixed: having/seeing an appropriate one is to constitute the experiential component in an account of the above form. So sense data must be such that for any of these one might have while seeing an apple, it may be conceded that you had it without prejudice to the question whether you saw an apple. This makes clear how Grice is treating seeing, and what genre his account belongs to. To use a term of John McDowell's, Grice treats seeing as a hybrid phenomenon, factoring into a purely mental, individualistic component and a worldly component. The mental component depends only on how things seem to the seer; its being the way it is depends only on his being the way he is, not on any extradermal state of affairs. The worldly component depends on how things are beyond the seer's skin, in ways that (on the view) are not, and cannot be guaranteed by how things are presented 'within'. (It is just this picture of inner and outer states of affairs that Austin and others attacked.)

Other members of the genre cover the expected cases: ones that admit of ringers. Such as knowledge. The idea would be: the way things are within cannot

guarantee the way the extradermal world is. So, for  $X$  to know  $P$  is for  $X$  to be in some internal state, which  $X$  might be in whether or not  $P$ , and for some further state of affairs to obtain—e.g.,  $P$ . This idea about knowledge has been attacked forcefully by McDowell (1982). There are similar ideas about thoughts: to think, e.g., that Grice was a philosopher is to represent things as one might with or without Grice (the inner component), while Grice exists and corresponds suitably to your representation (the worldly component). This idea about thought has been attacked by McDowell and by Gareth Evans (1982). There is, finally, a parallel idea about saying: if someone says, e.g., ‘Grice was a philosopher’, then those words are to be understood in a certain way, where that would be the proper understanding of them whether or not Grice existed. They then say that Grice was a philosopher in virtue of the additional worldly fact that Grice did exist and corresponds to the proper understanding of those words in the required way. The most famous attack on this idea is by Saul Kripke (1972). There is no reason for such attacks not to fit a hybrid account of seeing. They have been so applied by Paul Snowdon (1980).

Let us turn to Grice’s sanitizing efforts on behalf of sense data. The idea is: wherever I see something,  $Y$ , I have a particular experience,  $E$ . I need not have just the same experience every time I see  $Y$ . Neither need it be an experience every one has whenever he sees  $Y$ . (For these reasons alone,  $E$  cannot be seeing  $Y$ .) But  $E$  must be an experience of the right sort. To be so, it must meet these conditions: first,  $E$  is an experience I might have had whether or not I was seeing  $Y$ , or for that matter, anything; second, to have had  $E$  where some supplementary condition is fulfilled (namely, that  $E$  was properly caused) is for me to have seen  $Y$ . I need have had no *further* experience. Sense data are to play the role of  $E$ .

Is there anything to play this role? If so, what? Grice aims to say how to say what sense data are, so that there turn out to be some and they fill the role. Traditionally, sense data are introduced by an argument; they are what it proves there must be. As Grice recognizes, these arguments are invariably bad. So Grice proposes other means. He aims to show that we really have recognized such things all along; and to do that by tying sense data to ordinary locutions, which, he aims to show, describe something in particular, in each case of seeing, and which meet the above requirements on  $E$ .

Grice does not have a definitive proposal for tying ‘sense datum’ to ordinary language. He canvasses two. On one, the preferred locution is ‘ $X$  looks  $F$  to me’ (p. 226). The link is: wherever something looks  $F$  to me (I think so? It seems so?), I am having an  $F$ -ish sense datum (definition). The other (p. 237) links sense data to the ordinary locution, ‘It looks to  $X$  as if . . .’, the link being as before. Grice also suggests that one might cast a causal theory metalinguistically, restricting the term ‘sense datum’ to the compound ‘sense datum statement’.

It is at this point that Grice makes contact with those philosophers, such as Austin, who brought sense data into ill repute. It is because of the way he makes

contact that, from his point of view, it does not matter that the above proposals are neither specific nor precise. Austin suggests a number of different criticisms of proposals of the above sort. Grice is interested in just one. The sort that interests him is, for the first proposal: it is not invariably the case, whenever someone sees something, that there is such a thing as 'the way it looks to him', or any fact about the way it looks to him which holds *tout court*, independent of whether he is seeing anything. For the second: it is not invariably the case, where someone sees something, that it 'looks to him as if' anything at all. 'Looks as if' is not an expression with general application; it describes, and is true of, only very special cases.

The point is essentially the same for both proposals. Since it is more striking for the second, I confine attention to that. There is, though, one point about the first, of which I think Grice was aware. Suppose that the grass looks green to me (waiving all worries about whether that is odd to say). Then the *grass* looks green to me; a state of affairs that could hardly obtain without my seeing the grass.<sup>4</sup> My having a sense datum so defined is not without prejudice to my seeing the relevant thing. Such sense data do not fill the role of *E*. Nor are we familiar with anything that looks green to me which does not thus prejudice issues.

Grice would not cast his opponent's claims quite as above. Or if he did, he would quickly progress to the data on which he takes those claims to rest. The kind of claim that interests him is this. Suppose I am holding a perfectly ordinary apple before my face, preparing to take a bite. Then it is normally odd/wrong/a misuse of language to say, 'It looks to me as if there is an apple before me'. Against this, Grice deploys the full machinery of conversational implicature: it may be odd to say such a thing for any of many reasons, which need not be that it would be false. (Supposing that there *is* something thus said.)

As we now know, though, such remarks about oddity are likely peripheral to a more interesting point. Suppose that 'we would not say' 'It looks to me as if there is an apple before me' in such and such situations. There is likely a reason. First, there may be many distinct things to be said, by a person, at a time, in saying 'It looks to me as if there is an apple before me', many distinct contributions that 'looks to me as if' may make to wholes of which it is a part. What is said might vary, for example, with what things are thus said *not* to look like. However it is drawn, there is, e.g., the contrast between *that* (the way a pomegranate looks) looking to *me* like an apple (I don't know fruit), the fuzzy reddish thing (in fact a pomegranate) looking like an apple to my clouded eyes, that swarm of fruit flies making it look to me as if there is an (obscured) apple there, its looking to me as if there is an apple and not wax fruit, and so on.<sup>5</sup> Second, since there is thus a semantics-fixing task for facts of a speaking to perform, some

<sup>4</sup> I could not be hallucinating and having precisely *that* experience. Nor one of *anything* looking green to me; hence nor would I thus have a greenish sense datum.

<sup>5</sup> Changing the locution will change, but not reduce, this variety.

speaking may fail to perform it. Then nothing would have been said (to be so). More metaphysically, there is no *single* state of affairs that 'It looks to *X* as if *A*' describes; no particular experience it identifies. Nor, on most occasions, does 'what it looks to *X* as if' identify anything at all, much less some experience *X* would have had, *tout court*, whether or not he saw an apple. Moreover, as Austin pointed out, the differences between the things those words may say suggests that they are more like rendering a verdict than reporting an aspect of experience.

The above raises substantive issues on which the cogency of Grice's scheme turns. Grice does not get so far as seeing those thoughts as his actual opposition, though they are the ones that brought sense data into ill repute. They are not dismissed by remarking that certain things would be implicated if others *were* said; nor by showing certain suggestions to be cancellable, as, of course, they would be if *different* things are (literally) said in different speakings of given words with fixed meanings. Conversational implicature is not to the point here. The means Grice deploys to defend a supposedly sanitized notion of sense data are impotent.

### 3. Meaning

Grice's view of meaning covers three notions: an expression meaning such-and-such in its language (linguistic meaning); a speaker meaning such-and-such *by* his words (speaker meaning); and words having meant such-and-such on a speaking of them (utterance meaning). The first notion is meant to reduce to the second on roughly these lines: words *W* mean *X* (in language *L*) just in case *X* is what speakers of *L* 'standardly mean' by *W*. The second notion is meant to reduce to (other) intentions in roughly this way: for a speaker, *N*, to mean *X* by (his words) *E* is for *N* to 'intend to produce some effect in his audience' (a particular one, that is), and to intend to do this in a given way, roughly 'by means of the recognition of this intention'. The third notion plays this role: for words *W* to have meant *X* on a speaking is roughly for their speaker to have meant *X* by them (where there *is* an identifiable speaker). For non-ambiguous words *W* to mean *X* (e.g., in English) is roughly for them 'standardly' to mean *X* on speakings of them.

The quotes are from Grice's original 1957 article (1989, p. 220), which says little more as to what effect must be intended or the way it must be meant to come about. Grice says more about the intended effect in a 1968 piece (1989, essay 6). There he states that in 1957 he meant the intended effect to be, for a normal declarative utterance, that the audience should believe something, but that he now (1968) means the intended effect to be that the audience should 'think that the utterer believes something' (1989, p. 123). To think that the speaker believes something *is* to believe something. I take it Grice means: for the speaker to have meant (that) *X* by *E* is (1957 view) for the speaker to intend



the audience to believe that *X*, etc., and (1968 view) for the speaker to intend the audience to believe that the speaker believes that *X*. In the 1982 'Meaning Revisited' (1989, essay 18), Grice seems no longer to mind whether we think of things in the 1957 or the 1968 way. (See 1989, p. 300.) And the most important features of his view are in fact independent of such details.

Grice's project has several possible attractions. One is the promise of reduction of the semantic to something else. (To intentions first, perhaps as a way station to something less problematic.) Grice aims for analysis, but disclaims interest in what he calls *reductionism*, which he explains as the idea that 'semantic concepts are unsatisfactory or even unintelligible, unless they can be provided with interpretation in terms of some predetermined, privileged, and favoured array of concepts' (p. 351). He may, though, encourage reductionist interest when he says, e.g.,

'Meaning<sub>nn</sub>' ('non-natural meaning') does not look as if it names an original feature of items in the world (p. 350),

which, he says, 'supports the suitability of further analysis for the concept'. A second attraction (to some) is that Grice's analysis, applying, if at all, to much more than *linguistic* communication, promises to place linguistic acts within some wider domain. Third, Grice makes much of the idea of an intention such that its being manifest that someone has it guarantees that it is realized. Such intentions were not widely discussed before Grice. They seem interesting. Finally, there is a simple hankering after analysis, Grice's own stated motive.

### 1. Preliminaries

Grice's rough reductions posit non-existent conceptual links; they are not even roughly on the right track. The subject's complexity makes it difficult to see that. Four preliminary points will help.

*First.* Grice marks off an area in which his targets lie, 'non-natural meaning', by contrast with a sort of meaning outside his intended scope, 'natural meaning'. He makes it clear enough what natural meaning is. (Its most important feature is that it is roughly factive: 'These tracks mean that there are deer about' is right only if there *are* deer about; 'These tracks mean deer' only if deer are what (standardly) produce those tracks.) *Non-natural* meaning is defined by exclusion: anything meriting the term 'meaning' which is not natural meaning. Is this a unified domain? 'Mean' is idiosyncratically multifaceted (compared, e.g., with its Dutch or German cognates). If it seems just 'right' that linguistic meaning is, somehow, all a matter of speaker meaning, does it seem equally inevitable that *Bedeutung* is all a matter of *Meinung*, or *betekenis* all a matter of *bedoelingen*? If, as I think, not, that counsels caution. Further, even if what words mean is *somehow* all a question of what speakers mean, so of intentions, there are many

ways that might be. Grice assumes a naturalistic approach: linguistic meaning is to be some function of speaker's *actual* intentions. Intentions might play other roles.<sup>6</sup>

*Second.* 'Mean by', as it stands, fails to fit Grice's scheme. For whole utterances, and often for others, it typically points *away* from what words mean, and in no particular direction. This may be part of what Wittgenstein meant by:

The question 'what do I mean by that?' is one of the most misleading of expressions. In most cases one might answer: 'Nothing at all—I *say* . . .'. (1967, §4)

What *did* Wittgenstein mean by that?—a call for exposition of his point, *not* explanation of what his words meant. Similarly, 'Beliefs are not gold.' 'What do you mean by that?' 'They lack an essence.' 'You will regret it if you open that gate.' 'What do you mean by that?' 'The very large dog behind it will leap at your throat.' 'Our halcyon days are behind us.' 'What do you mean by that?' 'Old age is a drag.' / 'Western economies are on the verge of collapse.'

On the most natural use, what one means by one's words, if anything, is what one is driving at or suggesting. Hence Wittgenstein's remark: typically we do not mean anything by our words; we just speak them, and thereby say whatever they do say so spoken. What one may drive at in saying *X* is constrained in no specifiable way by, nor derivable from, what *X* mean in their language. An answer to a 'mean by'-question is *quite* different from an explication of what words mean, or even what one meant them to mean. Conversely, it would be amazing if there were *anything* English speakers *standardly* meant by most words, all the more if what those words mean in English were equatable with such a thing.

There are also cases where 'what do you mean by . . .?' asks for paraphrase, saying the same more comprehensibly. Rarely, if ever, does paraphrase explain what words *mean*. Consider:

'The only method . . . is for Mr. Todd to induce some second party to gather the actual experiences which Miss Rockmetteller wishes reported to her, and to convey these to him in the shape of a careful report, on which it would be possible for him, with the aid of his imagination, to base the suggested correspondence.'

'Could he put it a little clearer, Bertie', he said. 'I thought at the start it was going to make sense, but it kind of flickered.' . . .

'My dear old man, perfectly simple. I knew we could stand on Jeeves. All you've got to do is to get somebody to go round the town for you and take a few notes, and then you work the notes up into letters. That's it, isn't it, Jeeves?'

'Precisely, sir.' (Wodehouse 1988, p. 153)

<sup>6</sup> For example, *if* it is true that our understandings of words are our answers to the question what the speaker meant by them, or, better, how the speaker meant them, then perhaps what they meant, and what they said, is a matter of the intentions the speaker *ought to have been taken* to have.

Bertie said what Jeeves meant. Since Jeeves said what he meant, Bertie's reply also states what Jeeves *said*; rather than what he was driving at. It will hardly do as an account of what Jeeves's words mean, or then meant.

Cases still closer to Grice's purposes arise for proper parts of sentences: 'What (whom) did you mean by "Bill"/"that tall pointy building"?' 'exactly what do you mean by "reliable"/"colourful"/"heavyset"?' Occasionally, 'When you said "The walls are white", just what did you mean by "white"?' Again, answers to such questions seldom explain what the words used mean in English, or meant on that occasion.

Ambiguous words may sometimes fit Grice's purpose exactly: 'When you said "There's the vice-admiral", did you mean by that *the admiral in charge of vice*, or *the holder of the rank just below admiral*?' Those English words might mean either. To answer the question, choose, so state one thing they do mean in English. This, then, is one thing 'mean by' sometimes says. If that understanding of 'mean by' could be extended to all cases, then 'mean by' so used would serve Grice's purpose—if he could specify what understanding, or use, his 'mean by' was always to be taken to have, and do that without use of notions too close to what 'mean by' aims to analyze. (Intuitive grasp of what Grice has in mind would not be enough. In fact, I think we have that: by 'mean by' Grice means 'mean to say'.<sup>7</sup> But substitute that for 'mean by', and plausibility leaks from his analysis.)

*Third.* Grice persistently, though not consistently, uses 'means that' to describe what English sentences do, either as such or on occasions. (See, e.g., pp. 217 and 298.) As he uses it, this is a solecism.<sup>8</sup> 'Means that', in this connection, can refer *only* to natural meaning. 'Snow is white' (i.e., the fact that that is English) means that English is a Germanic language, all of its constituents being of Germanic origin. What 'Snow is white' does *not* mean is that snow is white (unless the whiteness of snow caused English to contain that sentence).

The above is no grammatical quirk. (I now assume the Wittgenstein–Austin view of language.) Suppose that a sentence non-naturally meant, or said, *that* such and such. The such and such should be possibly so or not. Some ways for the world to be should make it so (or not). Such a sentence would be rich in truth-involving properties. It would be true if the world is thus, false if the world is so; or true/false of those ways for the world to be. *Ceteris paribus*, it

<sup>7</sup> Note that what Grice means by 'mean by' is not a matter of his *meaning* 'mean by' to mean, or say, that. What he means is shown by the way he uses those words, and, if I am right, the only reasonable way of understanding them so used; the understanding that would make them so usable. If that is right, then our practice is against Grice.

<sup>8</sup> There is a related and equally grievous error. In speaking of inducing belief, Grice speaks of a speaker wanting a hearer to think '*p*' rather than that *p*. (See p. 300.) But, for much the same reasons that 'means that' is not fit for Grice's purposes, there is no thought in particular which is 'the one one would think in thinking "*p*"'—this is not a way of identifying something to be thought. So this way of speaking also fails to serve Grice's purpose.

should *be* true or, *casus quo*, false. That would make it unsuitable for saying the various things there are to be said on occasions in speaking of its topics. There are various things, e.g., that being white sometimes comes to. For some purposes, snow does not count as white. (Looked at close up, it is transparent, thus contrasting with milk or sheets.) Where that is what being white comes to, the appropriate truth should be expressible in saying ‘Snow is not white’. English appears to permit this. The price of that flexibility is that no English sentence is, as such, rich enough in truth-involving properties to mean, or say, *that* anything. Rather, English sentences vary their semantics, so what they say, and its truth condition, across speakings. Grice’s solecism depends on this not being so. Before converting solecism to neologism, that presumption must be argued.

Grice capitalizes on his solecism. He wants something, *X*, which is both what English speakers ‘standardly’ mean by a sentence, *S*, and what *S* means. For speakers to mean that just *is* for *S* to mean that. But what someone means by a sentence is, typically, *that* thus and so. Which is not eligible to be what a sentence means. The solecism papers over exactly that gap. Perhaps a sentence’s meaning is an *aspect* of what people mean by it when spoken. That aspect must still be factored out.

*Fourth.* A semantics for words must be coherent. By a semantics I mean a set of semantic properties. A coherent one is a set *some* words, and what they said, could have.<sup>9</sup> Not every semantics is coherent. Having a semantic property excludes having certain others. There is the property of saying something to be blancmange, and that of being true of my piano. My piano, on no view, counts as blancmange. So no words could have both those properties.<sup>10</sup>

Questions of coherence are not resolved by the opinions, nor by the intentions, of a speaker. I cannot get words of mine to have the above combination of properties by intending that they should, or that those properties should cohere. Rather, whether a semantics coheres is a matter of whether it is rightly, that is, most reasonably, so regarded. Which, in turn, is *roughly* a matter of what we are prepared to recognize. What *we* may say truly about the semantics of a speaker’s words is constrained by which semantics anyway count as coherent. Such is not to be read off the speaker’s state of mind.

The semantic coherence of intentions for words is seldom a problem. One reason is that semantic intentions are typically sparse (as *Zettel* §4 points out).

<sup>9</sup> Suppose I speak ambiguously, as in *double entendre*. Then there are two coherent semantics such that my words are properly understood to have said what the first, and *not* the second, fits, and *in addition* to have said what the second, and not the first, fits. They are not to be understood to have said what the union of the two semantics fits. If the union is incoherent, there is no such thing to be said. There is no explaining *double entendre* except in terms of coherent semantics; semantics for saying *something*.

<sup>10</sup> Words may, of course, state the inconsistent. They do that by having a coherent, and coherently specifiable, but unsatisfiable, condition for truth.

We typically mean our words to mean what they *do* mean, whatever that might be, and to have the additional semantics that words which meant that would have under the circumstances of our speaking. Typically, the proper assignment of semantics to words is the semantics those words have anyway in their language, supplemented in the most reasonable coherent way, given the functions those words would serve so produced. Given that the speaker's intentions were thus sparse, no further reference to them occurs in assigning semantics correctly.

Mutually incoherent speaker intentions can occur. Someone says, 'Give me a dace'. He means his 'dace' to mean *dace*, a small freshwater fish excellent for *quenelles*. But he may think that dace are a certain large saltwater fish, so also mean his 'dace' to speak of that saltwater fish. It is not possible to ask for dace and thereby to ask for a fish of the sort he is thinking of (though, of course, one can ask for one of each). His words would do that if they realized both intentions. They would thus have an incoherent semantics. They may, in fact, have had a coherent one, but, if so, only because their semantics need not be a direct translation of his intentions. (And cannot be if he asked for anything at all.) Surroundings may show, e.g., that one of his intentions is most reasonably treated as inconsequential.<sup>11</sup> If surroundings need to show this, and might, then that shows that the semantics words *have*, what they actually said and meant, is not a construct out of actual intentions.

## 2. Evaluation

I now turn to evaluating Grice's various proposals.

*Utterance meaning.* For sentence (linguistic) meaning, abjuring 'that'-clauses is mandatory. (See point three.) For utterance meaning, we may choose: retain the 'mean' and drop the 'that' from 'mean that'; or retain the 'that' and substitute 'say' for 'mean'. Accordingly, we isolate one of two distinct phenomena. (Retain both 'mean' and 'that' and we study *natural* meaning.)<sup>12</sup>

The first phenomenon arises for ambiguous expressions: the English 'Mary had a little lamb' may mean *ate a bit* or *kept one*; on a given speaking it may have meant exclusively the one or the other. The second phenomenon concerns saying things to be so, thus what may be true or false: Sam said, mendaciously, that Mary Jones (the professed vegetarian) ate some lamb last night. Grice's

<sup>11</sup> In many cases, but not all, the fishmonger who gave dace committed no fraud; he gave what was asked for.

<sup>12</sup> Note that the issue here is what (someone's) *words* meant. Grice gives examples like: 'In uttering "My trouble and strife is coming" he meant that his wife was coming'. That is about what the *speaker* meant. His *words* meant: *my wife is coming* ('trouble and strife' being Cockney for 'wife'), and *said that* his wife was coming. They *meant* that his wife was coming only if the impending presence of his wife regularly causes this man to blurt them out.

ambitions plainly exceed the confines of the first phenomenon. I thus take him to speak of the second, even though the resulting view of saying fits ill with his view of it when concerned with conversational implicature. I will comment on the tension later.

Given the provisional interpretation, Grice's proposal is: words said that  $X$  (roughly) just in case their speaker meant  $X$  by them. Now by point two, what people meant by their words typically *contrasts* with what was said in them. Pia *said* that coffee would keep her awake; she *meant* by that (but did not say) that she didn't want any.<sup>13</sup> We might correct that problem by replacing 'meant by' with 'meant to say', or 'meant his words to say', or stipulating, or isolating, a use of 'mean by' on which it means that. That would seem to capture the sort of equation Grice has in mind. But is this new equation at all plausible? Can success in realizing one's intentions always, in the case of saying, be *that* easy; guaranteed in principle? If so, why bother, in mundane affairs, to speak of what people *meant* to say, or etc., at all? (Not: 'What the candidate *meant* to say was . . .', but simply: 'What the candidate *said* was . . .'.) Or, again, is doing business that intrinsically risky? (The customer said 'dace'. The fishmonger agreed to supply it, and *did* supply dace. Could the mere fact of the customer's intention for his 'dace' put the fishmonger in breach of contract?)

Even if one swallows such results, there remains the problem of principle posed by point four. We may typically be credited with masses of diverse intentions for our words. Translate them slavishly, by naturalist formula, into our words' actual semantics, and you will often assign our words incoherent semantics. It is implausible that in *no* such case was something said. That would make saying things enormously difficult. But no words, nor anything said, can have an incoherent semantics. To read a semantics for your words out of your intentions, we *must*, in such cases, appeal to the reasonable construal of them. To do that is to appeal neither, anew, to your intentions, nor to an antecedently fixed function for mapping those intentions onto something.

The point holds even where all your intentions may be coherently realized. The semantics of any words is closed under requirement. So the rule is: if you intend your words to have semantic property  $F$ , and your intention is realized (you intend them to speak of dace, and they do), then they also have any other property that having  $F$  requires, and lack any that having  $F$  excludes. What does application of that rule yield for given words? That depends on which semantic properties are correctly judged to require or exclude which others. Correct judgement in such matter is no question of speaker intentions.

Suppose, now, that we replace Grice's 'meant that' with 'mean', so take him to be treating the first phenomenon. If English words, such as 'Mary had a little

<sup>13</sup> This problem does not arise for 'mean' *sans* 'that'. A speaker may have meant *riverside* by his 'bank', and that 'bank' may accordingly have meant that. There are, I think, other objections to equating these phenomena. These will arise later.

lamb', may mean one thing, and may mean another, then when it comes to making precisely that choice, we go out of our way to let the speaker's intention, where he has one, decide. (It is wrong to elevate this policy to Gricean analysis; there are limits to tolerance.) The relevant intention is that, of those familiar readings the words do have in English, such and such be the one they bear as spoken. The intention identifies its object by appeal to familiar facts of English, e.g., as 'the meat-eating sense'. Words cannot mean what they do in English in virtue of intentions of this sort. Wavering between meaning and saying, as Grice's solecism encourages, may yield a dual illusion: what words said looks closer than it is to what speakers meant to say (as close as the words' meaning as used is to what speakers meant them to mean); speaker's intentions as to what their words should mean look unduly close to something out of which what words *do* mean might be constructed.

Suppose that, rather than following Grice's path, we begin with the observation that there *are* sometimes facts about the semantics words had, and what they said. We might then ask when it would be correct to judge words *W* to have semantics *S*. Suppose, e.g., Sam judges Pia's 'The walls are white' to be made true by such and such way for walls to be. What would show Sam mistaken? Perhaps some fact Sam did not know about Pia's speaking or the way it fitted into its background. That fact might change his view of how the words were to be taken. Or, perhaps, Sam's reaction to that speaking is unreasonable. That is something reasonable language users, on the whole, could see, if the flaw were properly spelled out.<sup>14</sup>

Now for a non-Gricean hypothesis. If there is no such demonstration that Sam is mistaken, then he is not. He has understood the words as they were meant to be, that is, correctly, so assigned them semantics they have. Roughly, then, words *W* have semantics *S*—*inter alia*, said that *X*—if they would be so understood, where there is no demonstration on above lines that that it is a mistake. This is one alternative to Grice's account. It appeals to what reasonable, competent language users would recognize when deprived of no relevant fact. It has little promise as *analysis*. That need not be a criticism.

This alternative does not tend in a naturalist direction. It allows that speakers' intentions *may* have great importance in fixing the semantics words actually had, as they would wherever the reasonable and competent took them to. But it also allows such intentions to be weighed up differently in different cases, and sometimes to be trumped by other considerations. If it is true, as the Gricean insists, that our understandings of words always represent our answers to some question as to what the speaker meant, it is consistent with that that the correct understanding of words should be fixed by the answers such questions ought to have been taken to have (by the properly informed, in those circumstances). Nor need the conditions for being a correct understanding ever be specifiable in

<sup>14</sup> Obviously a few wrinkles are called for. For further discussion see Travis 1989, ch.2.

any terms from which such evaluative, normative judgements are wholly absent. (Such judgements may always be ones we, and other reasonable beings, are able to supply.)

There is a tension in Grice. Suppose one is committed to the picture of what is (literally) said with which Grice backs his account of implicature. That will incline one to discount any speaking-variable phenomenon with instances not derivable from what words mean. Such a phenomenon will not be seen as bearing on what was (really) *said*, or the condition for its (literal) truth; it will not be seen as genuinely semantic. Grice's speaker meaning is unconstrained by what words mean: *some* speaker might, on occasion, mean just about anything by anything. For a philosopher gripped by the picture just mentioned, it is tempting to regard any speaking-variable phenomenon thus floating free of (linguistic) meaning as a question of *mere* speaker meaning, and to see it as thereby demoted from the ranks of serious semantic issues; certainly issues of what was actually *said* on an occasion.

Treating what a speaker meant by his words as *mere*, so contrasting with what those words actually said, or stated, is in conflict with treating utterance meaning as even nearly *equivalent* to it, if utterance meaning is anything more than a choice between a stock of ambiguities fixed antecedently by English. Given the equation, to identify a phenomenon as one of speaker meaning is not to diminish its *semantic* import. Nor is that import diminished on the other way of purging solecism, by speaking of what words meant on an occasion (e.g., *ate* rather than *kept*). If Grice intends to identify some *other* phenomenon with the solecistic 'words *S* meant that', it is obscure what it, or its interest, could be. I suspect that Grice's solecistic shift from 'say that' to 'mean that' simply disguises the fact that he has two mutually inconsistent views of semantics and saying, each well suited to the separate purposes of different philosophic moments.

*Speaker meaning.* That meaning *X* by *Y* is a matter of having the right intentions is eminently plausible. English permits 'What do you intend by that?' to substitute for 'What do you mean . . .?', both questions being recognizably close to 'How do you intend/mean that?'. Dutch and German, which mark a cleft between linguistic meaning and speaker meaning, describe the latter by words which also describe intentions. But which intentions matter, and how?

Recall point two. 'Mean by' is closest to what Grice wants where ambiguity in a language links it to choice between some one specific reading for words and some others. It is next closest where linked to calls for paraphrase. In both cases, to ask the speaker what he meant is to invite a second try at saying what he meant to say.<sup>15</sup> Correct answers to such questions are provided fully by what such second tries would yield—that is, where those second tries succeed. There

<sup>15</sup> That is, was saying, where those words bear a progressive aspect. So read, it does not follow from the fact that the speaker was saying *X* that he actually said it.



are plausibly two requirements for success: first, such respeakings should count as, in some sense, saying the same as what the original might have; second, they should not raise anew similar questions as to what the speaker meant by them. The speaker will have *meant* by the original what he *said* in the respeaking. What he thus said should be evident.

By point four, the semantics of words must ultimately be fixed, not by speaker's intentions, but by what those, and any other relevant circumstances, make reasonable; it is an issue of which *coherent* semantics best fits those facts. So, too, a point comes where a speaker's respeakings cease to be informative; further efforts to say the same tell us no more. Something must count as the speaker's having found the right words; ones that fit his intentions. What *those* words say is just a matter of what such words, used in that situation, *would* say, that is, what they would most reasonably be taken to say. That is not a question for speaker's intentions to decide. So a speaker's intentions may tell us which words to treat as equivalent to his original ones as meant; but not what semantics such further words, so used, have.

We produce signs; where these are misunderstood, or we think so, we may correct that failure by producing others. Our meaning effects an exchange of our original signs for others. Conversely, it is fixed by those exchanges it would license or condone (or which we would where we counted as stating it).<sup>16</sup> No more than what would thus be fixed is involved in meaning words in one way rather than another. But such exchanges fix a semantics for our words only given a background against which what we exchange for counts as having this or that semantics already.

Where what one means by words is a matter of the semantics one intended for them, it is a matter of exactly that: what one intended them to say, or how one intended them to be (to be taken). Here Grice's conception of the business of philosophy—analysis—again forces his analysis along unnatural paths. For the question arises: why should to intend *X* be, in the nature of the concepts, to intend *Y*, *Y* distinct from *X*? Why should intending (one's words) to say *X* be just the same thing as intending anything *else*? It is implausible that it should be. (For one thing, if it were, our declarations of intention would be defeasible in ways they are not.) For Grice, though, some such equivalence *must* hold. He cannot stop with the kind of account I have just sketched. For, as he says in rejecting another approach, a

position hardly seems satisfactory when we see that it involves attributing to speakers an intention which is specified in terms of the very notion of meaning which is being analyzed (or in terms of a dangerously close relative of that notion). Circularity seems to be blatantly abroad. ('Retrospective Epilogue', p. 352)

<sup>16</sup> Of course, we also show how we meant words in the way we use and treat them. But it is even more obvious here that not every fact as to how we would treat our words can be directly reflected in a semantic property they in fact have. We cannot just will our words to have simply any semantics we might wish for them.

Relativized meaning is an intermediate analytic stage between nonrelativized meaning and a 'semantics-free' ('S-free') paraphrase of statements about meaning. So, given our present course, the fact (if it should be a fact) that there is no obviously available S-free paraphrase for . . . would be a reason against selecting . . . as an approved specimen of relativized meaning. (p. 353)

*Meaning to say that* will not do for the sort of way station towards eliminating semantic notions that meaning by was meant to be; for one could hardly say what *X* meant to say without invoking semantic notions.<sup>17</sup>

*Linguistic meaning.* Speakers typically mean their words to have the property of meaning, as used, whatever it is they do mean in the language spoken (perceived ambiguity aside). When we say 'lamb', we mean that to mean *lamb*, whatever that comes to. So, if 'lamb' means *X*, we may typically, though not always, be said to have intended our 'lamb' to mean *X*, and even to have meant *X* by it. That far from exhausts our intentions for our words' semantics, nor need all these intentions be consistent. But it is that aspect of intention which might be expected to remain constant across (standard) speakings of given words. That is one of two reasons for expecting what words mean to reappear in what speakers mean by them, or at least in how speakers meant them. The other is that typically, though not always, speakers *know* what their words mean. Knowing that, it is usually foolish to try to make them mean anything else.

These two considerations make it unsurprising if, as a not invariable rule,<sup>18</sup> if English words mean *X*, *X* is standardly part of the way English speakers mean them, especially given some give in 'standardly'. What the words mean may well reflect the most invariant feature of how people mean them. If such empirical generalizations are to be expected, that is not to say that *for* English words to mean *X* is *for* English speakers standardly to mean *X* by them; that the two notions are thus conceptually linked. Above, I suggested that for words to have had given semantics on a speaking is *not* for their speaker to have intended them to, or to have intended anything else; there is no fixed function from a speaker's intentions to his words' semantics, no simple reading off the one from the other. So Grice's 'utterance meaning', whatever it is, is *not* constructible out of speaker meaning. If what words mean in English is what they standardly mean as spoken in speaking it, then my suggestion, if correct, should transfer to links between linguistic meaning and speaker intentions: if the semantics of spoken words is not analysable in terms of actual speaker's intentions, then neither is the semantics of words of a language, such as English. If the one analysis forges the wrong conceptual links, then, equally, and for the same reasons, so does the other.

<sup>17</sup> Much the same analysis of Grice's difficulties has been proposed by McDowell (1980), who takes a view of speaker meaning not unlike the present one.

<sup>18</sup> What words mean is also a matter of their history, sometimes even where that is largely forgotten.

The facts that might lead one to expect links between speaker intentions and what words mean ought equally to lead one to expect links between what, e.g., English words mean and the way in which English speakers standardly understand speakings of them. So one might at least equally well say: English words  $W$  mean  $X$  just in case English speakers standardly take speakings of  $W$  to mean  $X$ . That is at least as plausible as Grice's proposal. This line of thought suggests a slight, but crucial, modification. On the line, what English words mean is a matter of the *actual* understandings of English speakers. In addition to actual understandings, we recognize a second, distinct notion. Pol understood Pia's 'Mary had a little lamb' as saying something of Mary's eating history—perhaps rightly, perhaps wrongly. In addition to the way he *did* understand the words, there is the matter of the way they *were* to be understood: their proper understanding. Perhaps they *actually*, understood correctly, said something about pet keeping. Given this distinct notion, one might also say: English words  $W$  mean  $X$  just in case standardly, on speakings, they are to be understood (are properly understood) as meaning  $X$ .

If this last equation is at least true, or on the right track, then the prospects for an analysis of linguistic meaning in terms of *actual* understandings are exactly as good as the prospects that proper understanding is a logical construct out of actual understandings; that the one is a determinate fixed function of the other. (Proper understandings might depend heavily, or even exclusively, on actual understandings without any such construction being possible.) The prospects for Grice's proposal for linguistic meaning can be no better than this.

This last proposal *of course* promises no *analysis* of meaning, no formula for replacing semantic notions by others. That is a defect only given that such analysis is possible, and that we demand it (neither of which follows from the mere fact, if it is one, that 'meaning is not an original feature of the world'). Conversely, if the proposal is true, that fact is reason for doubting that such analysis is a reasonable ambition.

Nor does this proposal promise to explain how languageless creatures could become linguistic ones. We know of no case where they do. Human beings just *are* linguistic creatures, as cats just are not. For linguistic beings, what their words might mean is fixed in part by the sort of linguistic creature they are. (One could not have the intentions we do for our words independent of any knowledge of what they mean, or might.)

#### 4. Conclusion

Analysis is not *a priori* misguided. *A posteriori* it may be judged by its fruits. Grice, having practised it brilliantly, gives us the means to do that. In many cases, analysis is more fruitfully abandoned.

If not analysis, what? Frege, and later Chomsky, provide one fruitful model. We begin with a mass of facts about words' proper understandings, *inter alia*, about what they mean. These are facts we are prepared to recognize. For the most part, they are not open to serious doubt. We might begin to approach them either *via* a theory of *meaning*, or *via* a theory, or account, of *meanings*.

In the first case, we assume an adequate grasp of *which* facts are to be accounted for.<sup>19</sup> We then ask: for each such fact, under what (otherwise specified) condition would it obtain? On what conditions would 'lamb' mean *lamb*, someone's 'Grice' refer to Grice, grammar *G* be the grammar of English, or of the language of community *C*, etc.? An answer to a range of such questions might be thought of as a recognition theory of its domain. It would enable one to recognize, e.g., cases of words saying that Mary had (ate) a little lamb, without reliance on intuitive understanding of them. That would be analysis in Grice's sense. For a suitably chosen domain, it would be an analysis of meaning, of one sort or another.

The second approach aims for a grasp of what the various things are that words might mean, or say, or etc. It would seek to identify those features by which one semantics for words might differ from another—e.g., the minimal distinguishing features by which words might differ in what they mean, or what they said. In doing that, it would uncover, and specify, those semantic features which may join to constitute an understanding of words, so a proper understanding of some; those semantic properties which may combine to form a semantics some words have. In the ideal case,<sup>20</sup> a theory with these aims would provide means for constructing a unique individuating specification of each item in its domain, for example, one for each distinct meaning English words might have. It would then be a generative theory of that domain as, e.g., Chomsky uses that term.

The characteristic method of the second approach, in exposing the facts it aims to treat, is the method of contrasting pairs. Beginning with words *W*, look for some words, *W\**, which share as many already recognized semantic features with *W* as possible, but still differ intuitively from *W* in their proper understanding. Try to see how like *W* words can be while still recognizably differing in their semantics. Success may lead to recognition of new aspects of what is said in words, new semantic features or properties which may constitute, in part, a proper understanding. This method was pioneered by Frege. For example, begin with an identity statement, then find another sharing all non-controversial semantic features—notably, its referring parts co-refer with their counterparts in the original—which still is to be understood differently. Success would exhibit

<sup>19</sup> It is remarkable that Grice thinks specifying *what* words said or meant a generally unproblematic task. That is a pervasive feature of his work, shaping his view of implicature, as well as his approach to meaning.

<sup>20</sup> I do not mean to suggest this as an ideal we should strive for. In general, it is probably unattainable. It is not always a good way of thinking about what one is doing in trying to gain a grasp of what understandings consist in, or how they vary. But it is a good starting point, an idea that begins us on fruitful paths.

a new semantic feature, that is, a new respect in which words  $W$  and  $W^*$  may differ in their proper understandings. (Frege called the feature he thought he found *Sinn*.)

The second approach is, obviously, the one I recommend for a start. This reflects my conviction that we students of language are not comfortably masters of the distinctions that need drawing in specifying understandings, so that there is still much for this method to reveal: about the ways we organize information (in using and understanding words), and about the relations between words and thought. There are, I suggest, at least three notable results so far.

First, what, e.g., English words mean radically underdetermines what there is to be said in them used *literally*, even ignoring such phenomena as indexicality, or the variety of individuals to which a name or description might refer. Any English predicate—‘is white’, for example—may, on a speaking, have any of many sets of truth-involving properties: ones fixing what, so spoken, it would be, or is, true of. Correspondingly, there are many distinguishable things it may say to be so of a thing, while used to say it to be white. It is not as if for each sentence of English there is the thought which it is its function in English to express. (Or, for ambiguous sentences, several such.) Rather, each sentence is a vehicle for saying any of indefinitely many things, each predicate usable, literally, for marking an indefinite range of distinctions the world *might* teach us need marking.

Second, if words’ meaning underdetermines what they say, then it is wrong to suppose that, if  $X$  spoke perfectly meaningful, well-formed words, then grasp of their meaning alone makes evident *what*  $X$  thus said, or even that it follows that there is anything  $X$ , or his words, said. Pick an unnatural enough home for words, and they may express no thought at all. (It is no news that that might sometimes happen when a referent is missing: ‘That chimpanzee looks tired’, I say, pointing into empty space. Section 1 illustrates subtler ways for this to happen—how, e.g., one may mistakenly believe himself to have said something about what someone knows.) (Cf. Wittgenstein, *Investigations*, §117.)

Third, if ‘is white’, spoken literally, may say any of various things of an object, that is due to a corresponding metaphysical fact: various states of affairs sometimes count, and sometimes don’t, as an item’s being white. That metaphysical point should hold as well for semantic properties as for others, with the same result for semantic predicates. The upshot would be that the semantics given words count as having may vary with the occasions for understanding them in one way or another. The point would hold as much for words on a speaking as for words in a language: *proper understanding* is an occasion-sensitive concept.

With an adequate grasp of which facts our domain consists of—of what meanings, or proper understandings, are like—we *might* aim for a recognition theory for that domain, perhaps on the model of Gricean analysis—a theory of meaning, or of (proper) understanding. As Chomsky pointed out long ago, success at the project I have put first *may* kill one’s interest in this second

one. In Chomsky's case, the recognition theories he aimed to postpone came in the form of discovery procedures for grammars. These would say, in other terms, and in complete generality, when  $G$  is the grammar of the language of community  $C$ . Those who sought such procedures were guided by a conception of grammar which made it necessarily very simple, just as Grice's view of what is said makes that a simple affair. On examination, the complexity of human grammar shows the search for such procedures to be uninteresting. Similarly the complexity of understandings, on current evidence, offers little hope for a theory of understanding, or of meaning, of the sort Gricean analysis would be. First, as with grammar, the semantics we are able to perceive is a function of the organisms we are. It need not be what other organisms could construct out of facts of another sort. Second, proper understandings are the most reasonable ones. Why think we could capture reasonableness in the nets of analysis?<sup>21</sup>

<sup>21</sup> From the start, I have profited enormously from criticism and encouragement by John Campbell. This piece has several considerable blunders less for that. I am also grateful to the editor and the assistant editor of *Mind* for many detailed and helpful suggestions.

# 3

## Meaning's Role in Truth

### I

What is truth?<sup>1</sup> A notorious question, tempting many, for millennia, to dismiss it. But even if the answer cannot be a definition, there is much to say about what truth is.<sup>2</sup> Part of what truth is, one might think, is the way the truth of words depends on what they mean. But there is a widespread view, tracing back at least to Frege, on which we may say all there is to say about what *truth* is without so much as mentioning words. I hope to show this view wrong. Here is my plan. What words mean plays a role in fixing when they would be true; but not an exhaustive one. Meaning leaves room for variation in truth conditions from one speaking to another. What that non-exhaustive role is depends on what it is to have said what is true. Identify the aspect of truth which fixes this role, and the widespread view collapses. So, step by step, I will argue.

Paul Horwich (1990, pp. 4–6) answers the question ‘What is truth?’ with a version of the widespread view:

The proposition *that quarks really exist* is true if and only if quarks really exist, the proposition *that lying is bad* is true if and only if lying is bad, . . . and so on; *but nothing more about truth need be assumed*. The entire conceptual and theoretical role of truth may be explained on this basis.

That is certainly *one* way with notoriety. What exactly is Horwich’s answer? First, he evidently supposes that when he said ‘The proposition that quarks really exist is true iff quarks really exist’ he stated a condition for the truth of *something*, or at least specified a condition under which something or other would be true (and, moreover, identified what that something is). Second, with his ‘and so on’, he suggests that that is just an example of something; and that he could easily have produced examples other than the two he did. I take that suggestion this way.

<sup>1</sup> It would be fair to view this essay as no more than a working out of some ideas of J. L. Austin (whom this question might call to mind). I have in mind particularly his ‘Truth’ (1979/1950) and ‘How to Talk’ (1979/1952). The essay in its present shape evolved with a great deal of very patient help from my colleague Peter Sullivan, and from James Hopkins, Michael Martin, Barry Smith, and Joan Weiner.

<sup>2</sup> This also seems to be Frege’s view. For, while he held truth to be indefinable, he also said: ‘The meaning of the word “true” is spelled out in the laws of truth’ (Frege 1918–19, trans. p. 2).

Suppose 'S' is a sentence in which Horwich might then have stated something, or at least might then have spoken of something (in particular) being so. Then instead of saying what he did, he might have said 'The proposition that S is true iff S'. That would be an example as good as either of the two he in fact gave. Let us call any fact anyone might ever state in words of this form an *H-equivalence*. Now consider the corpus of all the H-equivalences Horwich either stated, or might have, had he chosen different examples. I take Horwich to claim two things about this corpus: first, that it captures all that it is for a proposition to be true; second, that when we have said what it is for a proposition to be true, we have (as good as) said all that need be said as to what truth is, full stop.

There is an Aristotelian platitude which, paraphrasing freely, is to this effect: to state truth is to say things to be some way or other, and, in doing so, to say things to be none other than the way they are. That fits speakers as well as words. There is a parallel platitude about propositions: a proposition is true just in case things are as they are according to it. Must Horwich deny such platitudes, or reject them as somehow misbegotten? Nothing so far requires this. He need only hold that such generalizations hold precisely in case his collection of H-equivalences do. Since Horwich need not reject such platitudes, I will not hold him to doing so. So reading him, I will call his answer to the question 'What is truth?' *deflationism*.

Horwich has said nothing so far about what it would be for words to be true, or to state truth: nothing about conditions for their truth, or conditions under which they would have stated truth, or about just what they would have achieved in being true.<sup>3</sup> At first blush that seems a lacuna. But there is a simple, and widely accepted, story on which that lacuna, such as it is, is not one in an account of what *truth* is. On the story, for words to be true is (nothing other than) for them to express a proposition that is true. (Let us not haggle about whether that says more about truth than deflationism allows.) Once it has been determined which, if any, proposition words expressed (or would express), it has been determined when they would be true: when that proposition would be true. If they were not true under precisely *those* conditions, they would not have expressed that proposition. To see when that proposition would be true (the thought is), just consult the relevant H-equivalence.<sup>4</sup> Moreover, whatever determines which proposition words expressed, that does not depend on what it is for words, or anything, to be *true*; or on no more about this than has been mentioned.

<sup>3</sup> Horwich does say that from any declarative sentence 'S' we may form a nominalization, 'The proposition that S', and then 'denominalize' it back into a new sentence. 'The proposition that S is true', where that is, in some sense, equivalent to the original 'S'. But that tells us nothing about when anything would be true. Nor do we learn more about when words would be true by backing that remark with Horwich's H-equivalences.

<sup>4</sup> Could any deflationist account of truth provide all the H-equivalences that might be needed for this purpose? An interesting question, but not one that I will pursue here.



Confidence in this last assertion is generally fortified by a view of what does determine which proposition words did, or would, express: bracketing ambiguities words may have in their language,<sup>5</sup> and modulo the objects (people, etc.) words spoke of on an occasion, which proposition, if any, they expressed is determined by what they mean. Part of their meaning what they do is that (modulo referents) there is a particular proposition which is the one they would express.<sup>6</sup> The words 'Pigs grunt', e.g., meaning what they do, express a certain proposition which we may refer to as 'the proposition that pigs grunt'. The words 'Fred is fat', spoken on an occasion, expressed the proposition, of a particular person who they called 'Fred', that that person was, at the time of their speaking, fat. And so on. On this view, if we know what words mean, we can always identify the proposition they expressed on an occasion, though if we don't know their referents, our identification might not be in a very helpful form.

I will try to show that this is a wrong view of what meaning does, and that, given the right view, deflationism is a mistake. Not that what words mean is irrelevant to when they would be true. Meaning fixes *something* words would do (and say) wherever spoken meaning what they do; something they are for, so also something about what they ought to do. Truth requires that they do all that sufficiently well, that is, up to the standard truth imposes. But all that meaning fixes allows for words to state truth, but also falsehood, of given items in given conditions. What meaning fixes often enough leaves both possibilities open. This means, I will argue, that these requirements for truth cannot be captured in the form '*If* words expressed the proposition *P*, then they are true only where the condition for the truth of *P* is satisfied'. A given proposition is true just where the world is thus and so (or so the deflationist picture asks us to suppose). But there is no *one* way the world must be to supply what is required for the truth of words with given meaning. On the contrary, for different speakings of words alike in meaning, there are different ways the world must be. To capture this sensitivity of truth's demands to speakings, we must see how those demands make the way the world is matter differently to the truth of different such speakings. To do that we must depart from the forms deflationism allows.

On one model, words, meaning what they do, say, or said, thus and so, and then are simply true, or not, according as *that* is so or not; no further factors need be mentioned in the story. On the model I recommend, truth depends on what words mean, the way the world is, and further factors: aspects of the circumstances in which words were produced. Whether, and how, any given such factor matters to the truth of particular words must depend on what truth is; on the particular way speaking it requires words to relate to those factors (or vice versa). Further factors can matter at all, and then as they do, only if truth is

<sup>5</sup> Ambiguous English words express the proposition they would express on one of their readings in English. Similarly for other languages.

<sup>6</sup> How it is that they mean what they do is assumed not to be a relevant issue.

a notion which demands, case by case, precisely those factors to be arranged in just those ways. Deflationism cannot allow truth to make such demands.

## II

How does what words mean relate to what is said in speaking them? Consider the sentence 'The ball is round', and two cases of its use. *Case A*: What shape do squash balls assume on rebound? Pia hits a decent stroke; Jones watches. 'The ball is round,' she says at the crucial moment. Wrong. It has deformed into an ovoid. Jones did not say the ball to be as it was, so spoke falsely. *Case B*: Fiona has never seen squash played. From her present vantage point the ball seems a constant blur. 'What shape is that ball?,' she asks. 'The ball is round,' Alf replies; truly, since it is the sort of ball a squash ball (and this one) is. It is not, e.g., like a very small rugby ball.

So there are both true things and false things—thus a variety of different things—to be said of a given ball, and of the way it is at a given time, in the words 'The ball is round', used so as to have meant (as used) what they mean (in English). What those words mean, or their meaning what they do, makes no one of these things 'that which those words say' (in English). What those words mean leaves it open for them to be used (in suitable circumstances) to say any of various things, each true under, and on, different conditions. There is no *one* set of conditions under which those English words, spoken of a given ball and time, would be, or say what is, true. Nor even one condition which is the condition *for* them to be true. If differences in truth condition make for different propositions, then what those words mean makes no one proposition the one (modulo referents) they express.<sup>7</sup>

The example is not special. Most English sentences behave the same. Supplied with referents, there is still no one fact, or falsehood, they would state. That is not what meaning makes them do. Though it might be nice to make a fuller case, I will not pause to do so here.<sup>8</sup>

We can see why there should be this variety of things to say if we ask what meaning does do. The words 'is round', in meaning what they do, speak of being round. In fact, I suggest, for them to speak of that is just for them to mean what they do.<sup>9</sup> For English words to speak of being round comes to just this:<sup>10</sup> if you

<sup>7</sup> If differences in truth conditions do *not* make for different propositions, then you can't treat the truth of propositions in the way Horwich does.

<sup>8</sup> See, however, Ch. 4 below.

<sup>9</sup> Other words, of course, in speaking of being round, might do so in a derogatory or laudatory way, and so on. That might be part of their meaning what they do. I assume nothing like that is so of 'is round'.

<sup>10</sup> Perhaps it is still worth saying: it is not for there to be some particular set of things which are 'those things the words are true of'. We have already seen reason enough why there can be no such set.

use them as meaning what they do, you will thereby speak of being round. At least, on any occasion of your so speaking, that is something you would then be doing. So if you want to speak of being round, e.g., so as to call something round, or describe it as round, a way of achieving your aim in speaking normal English is to speak the words 'is round' (in a suitable construction).

We may put this by saying: there is something the English words 'is round' are for; something they are for doing in speaking English. They are for speaking of being round; in suitable constructions, describing things as round, calling things round, etc. That describes something they *would* do, spoken meaning what they do. There is another way in which those words are for something (in speaking English). At first approximation, they are for describing round things. Roughly that is what they *ought* to be used for; used otherwise they would not state truth.<sup>11</sup>

If this is what meaning does, we can see why there should be contrasts of the sort with which this section began—both true and false things to be said, e.g., of a given item in using 'is round' to mean what it does, in calling the item round. The reason is that the concept of being round does not by itself settle how an object must be to be correctly describable as round. The concept of shape, e.g., does not decide under just what conditions an object must assume a certain shape in order for that to be *its* shape. If what 'is round' means settled such issues where the concept does not, those words would not express it, and so would not mean what they do. There are various reasonable views of what it is for a ball to be a given shape, e.g., round. Words which say it to be round may be rightly understood as doing so on one or another of these views. That makes for various things to be said in so speaking. So if we are to judge whether words were used for what they ought to have been used for—to describe the things they are, or were then, for describing (things rightly and truly described as round), we will have to look beyond merely what the words used mean, on the one hand, and the way the described item in fact is, on the other. That raises a question: to what else must we look? I suggest that to answer that question we must look to what truth is.

The occasion-sensitivity of what words with given meanings say—the plasticity, as it were, of meaning—might well make one wonder just what H-equivalences might say, and what they could possibly say that would make some collection of them exhaust what truth is. But I will not pursue those questions here. Rather, I will stick to the question what it is for words to state truth, and how their doing so depends on what they mean.

### III

Call words which, rightly understood, say, or aim to say, how things are *statements*. A statement purports to satisfy a certain general condition: however it says things

<sup>11</sup> To miss the way in which this is only rough is to obliterate just the features of truth I mean to highlight here.

to be, whatever it says to be so, things are that way. Meeting that standard is a condition on its truth. If we ask when a statement would meet the standard, or what it must do to accomplish this, one move would be to break an answer down into particular cases: words which would describe a thing as round should be used in conformity with the rule: use them only to describe round things; words which would describe a fireplace as having a fire in it should be used in conformity with the rule: use them only to describe fireplaces with fires in them; and so on. If, in making a statement, you conform to all such applicable rules, your statement will satisfy the general condition.

One hitch in using this answer *may* come in deciding which such particular rules do apply, in the circumstances, to your chosen words. But leave that worry aside. What would it be to have conformed to a particular such rule on a particular occasion—to have done what the rule then demanded? For example, you described a ball as round. On one understanding of its being round, it is so; on another, it is not. So was the relevant rule (properly understood) conformed to? What decides this? There are three possible replies.

The first is this. Given the way the ball is, nothing about what it is to be round determines, by itself, whether the rule was followed; for nothing about that makes the ball either count as round, or count as not. So *nothing* decides whether it was followed. So nothing decides either that such words met the general condition or that they did not. There is just no fact of the matter.

If this line of thought is right, it would apply equally to any instance of 'The ball is round', spoken of a ball like the ball in the present case. So no such instance would be a case of speaking truth (or falsehood).<sup>12</sup> But that conflicts with the fact that there are various things to be said of a ball in describing it as round; in the case of a ball like the present one, some true and some false. So this response must be rejected.

A second reply is this. Depending on how you view the matter of (a ball's) being round, you might count that ball as (then) round, or you might not. If you do, then count the rule as followed. If you don't, then count it as not. Count the words in question as satisfying the general condition, or as not, accordingly. In so doing, you will be judging the demands of truth just as they ought to be judged.

This cannot be right either. For one might perfectly well sometimes view being round in one way and sometimes in another. There are occasions on, or purposes for, which I would happily describe the ball as round, and others on which I would

<sup>12</sup> Just in case this seems not really awful, note that the same sort of problem arises again for most ordinary squash balls at rest. Look carefully: there is probably an indentation where the dot is. And suppose squash balls were made with a seam. Again, such things sometimes count as being round, sometimes not. As Descartes pointed out, in our sublunary world, whether something satisfies the criteria for being round always depends on what you are going to count as satisfying them. (I will not pause to argue with people who take the heroic view that that just means that no one ever speaks truth in calling something round.)

equally happily say that it wasn't. (If you are too inflexible ever to depart from one view of what it would be to be round, there are still always contrasts between you and other, perhaps equally inflexible, but also equally reasonable, people who take a different view.) The second response would, accordingly, commit us to count a given statement sometimes as true, other times as not. We would need also to allow that some people who held the statement true, and others who held it false were equally, and perfectly, right, so in no genuine conflict.

The notion of truth does not work like that. Perhaps being true is an occasion-sensitive notion in the same sense that being round is. So there might be a statement which sometimes counted as true and sometimes did not. I might sometimes speak truth, and other times speak falsehood, in calling given words true. But truth is surely not occasion-sensitive in the way the second response demands. One might consistently hold that what *Fiona* said in calling the ball round is true, while what *Pia* said in calling it round is not, all the while with one's own view of what one would then call round oneself. That is the whole point about occasion-sensitivity. That is to say that I might judge words which described the ball as round to do (or fail to do) what *they* need to do for accuracy independent of whether *I*, at that particular moment, would count that ball as (then) round. I need not change my view of *Pia's* words simply on the basis of how I myself, at that moment, would describe the ball. Similarly, if Jones judges *Pia's* statement true, while *Fiona* judges it false, Jones and *Fiona* are not equally right—or right at all—merely on the grounds of the view each takes, at the moment, of what being round would be. Again, this cannot be how accuracy, in the present sense, is to be assessed.

Only one option remains. Whether words were spoken in conformity with a specific particular rule, so whether with all applicable ones—thus whether they count as satisfying the general condition—depends on the speaking in question, and not just on the state of things spoken of. Given the way the ball is, some words which described it as round may count as satisfying the condition, while others which so described it do not. What would count as satisfying it—the way the ball must be—is relative to the speaking, or describing, being assessed. Truth, that is, requires different things of different such describings.

To assess correctly the truth of words which described a ball as round, one needs to take the right view of what being round would come to—one view among many possible ones. The general condition, and the relevant particular rule, yield results only on an understanding, and the right results only on the right understanding. But what view is right for this purpose? That is a matter of the sorts of merits truth involves. For if, as we have seen, the circumstances of the speaking matter to what truth requires, something about what truth is, is needed to make such circumstances matter in one way rather than another, so in the way they do. So there must be something about what truth is which fills that role. Truth is sensitive to circumstances; so there is a way words must relate to their circumstances in order to be true.

We have reached a major change of perspective. The first two responses saw truth as the product of exactly two factors: first, which descriptions words gave of things—how they described things as being; second, the way those things in fact are, or were. From this perspective, a judgement of truth is simply one as to the match between these factors. Truth, from our present vantage point, is a product of three factors: the above two *and* the character of the occasion on which given words were used so to describe things. For a statement to be true, it must relate in the right way to a pair of factors: the way the world it describes is; and the circumstances of its making. So a substantial question arises. How must the statement relate to the second of these given the first (the way the world is)? Or, the same question put differently, how does the second (the circumstances) matter to how the statement must relate to the first (the way the world described is)? The answer must depend on something about just what sort of an achievement truth is. Without relevant substance to that notion, there *is* no answer.

The relevant aspects of what truth is, whatever they are, are not fixed by H-equivalences. For those, at best, relate but two factors: descriptions (individuated somehow) for words to give of things; and the world thus described. They are silent on how the circumstances of giving a description matter to its truth; that is, on the relation about which we are now inquiring.

Meaning, while it does not decide when words would be true, does not do nothing either. What words mean imposes definite conditions on their truth, but ones leaving a residual question: what would it be for those conditions to be satisfied *in the right way*—in the way, that is, that truth, in a given case, demands. It does not contribute to an answer to insist that if words say things to be *thus*, then things must be *thus*.

Words which satisfy the general condition, so say things to be only as they are, are *ipso facto* true. That does not mean that that condition is the only one truth imposes. If it were, what truth required could not choose between standards for assessing satisfaction of that condition, as we have seen it does and must. Rather, given what words mean, so how they describe things, they must satisfy the condition in a certain way: the way things are must count as things fitting those descriptions by those standards in achieving which words would relate to the occasion of their speaking as they must in order to be true. That selects between standards only insofar as there is a way truth demands that words relate to their occasions.

#### IV

There are two ways of thinking of the result just reached. To approach the first, consider the words 'Mary had a little lamb', a means in English for speaking of any of several things—eating ovine, or keeping an ovine, for example. English

provides a definite stock of readings for these words. On a given occurrence they may bear one or another of these. We might try to use readings as a model of ways of understanding words, thinking of a stock of understandings there are for words to bear. There is, e.g., a stock of understandings of 'The ball is round' for each of many understandings of what it would be to be so shaped. One reading words have in English may allow for any of many understandings of a speaking of them. But just as choosing the right reading fixes what words meant as used, so choosing the right understanding fixes what they said. For words which said something in particular, there is always some one understanding in the stock which is the one they bore.

One might further think: which understanding words bear depends on the circumstances of their speaking; when things would be as said to be on a given understanding does not. Understandings, so conceived, extract content from circumstances. Circumstances play no further role in determining conditions for truth. Deflationism, and its use of 'proposition', depend on exactly that idea.

Though I think this a bad picture, this is not the occasion to argue the point. Given the picture, the present result takes the form: truth is a determinant of content. Which understanding given words bore is fixed, in part, by how, for them to be true, the description they gave must mesh with the circumstances of its giving.

But here is another picture. Understanding requires sensitivity. Understanding words consists, in part, in sensitivity to how they fit with the circumstances of their speaking. Part of that is sensitivity to how they need to fit in order to be true. So adequate sensitivity requires grasping what truth is, and how that notion applies in particular cases. The sensitivity to words that grasp permits is, to coin a phrase, an *Auffassung* which is not a *Deutung*, and may constitute one's understanding of them.

Understanding a statement, we may agree, entails grasping well enough how it, and specifically its truth, are to be assessed. One might suppose that such a grasp, or what follows from it, could be made entirely explicit by some specification of an understanding there is anyway for words to bear, and by what follows from bearing that one (e.g., that they expressed 'the proposition thus and so'). But perhaps, as Wittgenstein suggested, a proper appreciation of words is not exhaustible in that way. If we do not assume so, then we may see truth, not as selecting from some stock of items each of which relates to the world in an entirely fixed way, but rather as guiding the treatment of words with given content in given circumstances, forming our perceptions of when to be satisfied with the descriptions they gave.

## V

The main result of this essay has been reached. What words mean imposes a condition on their saying, on a speaking, what is so. Different occasions

impose different standards for satisfying that condition. Something about what truth is makes occasions matter to such standards. Deflationism cannot recognize such elements in truth. This section explores, tentatively, what they might be.

What truth demands of a description is properly measured by different standards for different givings of it. What fixes the right standards for judging given words? One way to see what might do that work is by looking, not at anything anyone actually said, or might say, but rather just at descriptions to be given, and situations in which they might be. What, in given circumstances, would a given description describe (be for describing)? What is then so describable? Perhaps different circumstances for describing things yield different answers. To see whether that is so, we need to look at pairs of a description and a thing to be described.

For relevant pairs we might consider 'The ball is round'<sup>13</sup> and a ball just off the racquet, or 'Mary ate a little lamb' and Mary's meal of *mioleira*. Or, again, 'She is locked in' and Pia in a room with the doors locked, the window open, and a long drop; 'There's a fire in the fireplace' and a fireplace into which a lit ball of paper has been tossed; 'The oven is hot' and an oven at 140°C; 'The cat is puce' and a cat dipped in puce dye.

Let us begin with a lemma. One would not count a 140°C oven as hot for purposes of baking pizza. A hot oven, for that purpose, is much hotter. By contrast, for purposes of removing a rack bare-handed a 140° oven would count as hot. Something may be describable in a given way for some purposes, but not for others. There is such a thing as describability for a purpose. That is the lemma.

Why such facts, and why the particular facts there are? In one case the answer could go like this. One scheme for classifying ovens as hot or not reserves 'hot' for the upper part of a normal oven's working range. Ovens being as they are, that makes (say) 220° a fair dividing line. Nothing makes that scheme right, wrong, or implausible in any absolute sense. Choose it and the description 'hot' has a clear use in pizza baking: if an oven is so describable, the pizza may go in. For that purpose, no other scheme recommends itself (to us) so highly; which is enough for that scheme to decide, for that purpose, what is hot.

The culinary facts here are familiar. Familiarity breeds expectations. If an oven is described as hot, and we take that to be for purposes of pizza baking, *ceteris paribus* we will expect the information we are thus given to be usable as just described.<sup>14</sup> We will expect this in two senses. First, if we trust those words, we will be surprised if the oven is not pizza-ready. Second, we will (normally) expect that much of the words if we are not to count ourselves as misinformed by (in)

<sup>13</sup> For brevity, I quote words to indicate the description they give.

<sup>14</sup> Virtually every sentence, here and henceforth, describing how we would perceive words, calls for a *ceteris paribus*. Rather than writing it, let us take it as read.



them, so count them as stating what is incorrect. Our expectations thus form our perception of just what was said in giving that description.

We (often) perceive a description, as given, as for some given purpose(s), where, for those purposes there are uses that its fitting ought to have; uses for the information it ought to have provided. We count the description as giving correct information only where it did what it should have done to serve the relevant purpose(s). Such perceptions thus exhibit two ingredients—a purpose and uses words would have in serving it—which fix a standard for the truth of the words perceived.

We can move from the lemma to facts about what descriptions would describe on an occasion if we can find occasions on which certain purposes are the ones a given description ought to serve, in the sense that it would be right to expect that much of them. Such is certainly part of our perception of some occasions. We are, e.g., plainly engaged in making pizza. We would expect a description of an oven as hot to be reserved for that purpose. We would, correspondingly, perceive it as one that would give correct information only if it served that purpose. (We do sometimes make judgements as to what would, and would not, count as hot. These are conspicuously not judgements as to what is hot on such-and-such understanding of being hot. They are non-trivial in ways that the latter sort of judgement is not. Nor are these judgements as to what one could ever correctly call hot. Reminded that they could not be that if correct, we freely appeal to purposes for which something is being called hot; ones we initially took as read.)

It would be useless to be told that the oven is hot if we had no idea of the standard by which that was to be judged. Our perceptions of purposes to be served, and uses words would have in serving them, provide just such a standard. They thus allow descriptions of things to be the sort of good to us they are. These perceptions of occasions are perceptions of what it would be, on them, for a given description to describe *truly*, or for words which give it to state truth; to provide information which is correct. Their structure thus reveals some ingredients in truth, or what we are prepared to recognize about it. Part of the idea of truth is that a description (of something), to be true, must satisfy a general condition different in kind from conditions to the effect that what is described as *thus* must be as thus described: it must serve all the purposes that must be served (for truth) on that occasion, by having all the uses it ought in serving them. Part of this idea is that, for a description, and an occasion (on which there are facts as to what that description would describe truly), there are definite purposes truth demands be served, and uses which truth demands the description have in serving them. Another part of the idea consists in particular perceptions of how the concept of truth would apply in particular cases; what ought to count as, or be demanded for, describing truly on the concrete occasions on which the description might be offered—which purposes in particular, the description, so offered, would have to serve if true.

Our perceptions of descriptions on occasions thus link truth with use, that is, when words would state truth with the uses they ought to have to be true. Such is our idea of what truth, or stating it, is. Without that idea, it would be indeterminate when words had said things to be as they are. The idea of truth, we have seen, cannot exhaust itself in a cloud of ideas to the effect that words which say things to be *so* are true only if things *are* so.

## VI

So far we have not considered how to treat anything anyone actually says. We have only looked at what descriptions *would* be for on occasions. If, in that way, we can identify ingredients in truth—demands that words serve purposes, be appropriately useful—then, from one point of view, *tant mieux*. But we began by asking how to judge the truth of what is said. I suggest that the standards for that are fixed by just those ingredients in truth which we have so far identified. To see that we must see past a complicating factor.

To begin, our perceptions of a speaker's words, where we take ourselves to understand them, contain the same ingredients as our perceptions of descriptions on occasions. We see words as taking responsibility for serving certain purposes, in that we will count them as having said what is correct, so true, only where we count these purposes as (adequately) served. Jones says. 'The oven is hot'; Pia inserts the pizza. On learning, shortly, that the oven is at 140°, she will take herself to have been misinformed. On her understanding of the words, what they said is incorrect. Such are the sorts of understandings we take words to bear. Again, without such perceptions, we would have no standards by which to judge the accuracy of what is said; so there would be no judging it.

Pia may perceive wrongly. She may not grasp what the circumstances of the speaking are, or may miss some crucial fact as to what one could do if the oven were, in the right sense, hot (perhaps it is a trick oven), or may be plain unreasonable. Sometimes, though, we do no such thing. Sometimes we understand words. If Pia did no such thing, then her perception is correct.

What would Pia have to get right about the circumstances for her perception of Jones's words to be right? In the normal case, I suggest, the purposes someone's words should serve to be true are just the purposes the words that person used should serve, in the circumstances in which he spoke, to be true. So what Pia would have to perceive is what the words Jones used (the English 'the oven is hot') would be for on that occasion. But *so* much importance has been attached in the last forty years to a speaker's intentions, and there has been such a strong temptation to suppose that a speaker may fix the standards by which the accuracy of her words is to be judged merely by intending that they express 'the proposition

(such-and-such)', or by making this evident enough, that it is perhaps worth distinguishing two sorts of case: a default case and exceptional ones. I will explain that idea with an analogy.

In the normal course of events, in speaking English, we mean our words to mean what they do mean, and rely heavily, and *rightly*, on that fact, and on the facts as to what words mean in English, for fixing what they did mean (and say) as used by us. Jones asks, 'Could I have a brill?', meaning her 'brill' to mean *brill*. *Ceteris paribus*, that is what it did mean, which means that, in so speaking, she asked for a brill. If you want to know what a brill is, consult a dictionary. For further details, consult an ichthyologist or chef. In matters of what words meant as used, that is the default case.

Exceptionally, Jones may have special ideas about what a brill is. She may think, e.g., that 'brill' means bream. Moreover, she may make it clear that she means her 'brill' to speak of bream. Perhaps in some such case 'brill' as she used it did mean bream, so that what she asked for was a bream, and not a brill. Suppose so. Then there is a default case—what her 'brill' meant in the absence of special reasons provided by her for thinking otherwise; and there is the exceptional case where she supplied such special reasons. It remains so that in the default case what we have for working out what her 'brill' meant is just what 'brill' means. That is often just the way the speaker, and everyone else, wanted things to be.

Similarly, in the normal course of events we just talk, meaning to use descriptions for what they are then for—for what is then so describable—and taking ourselves to be doing just that. We rely on the facts as to what our descriptions *would* then describe for fixing when things would be as we described them—though, unlike fixing what our words meant, the relevant facts are not to be found in dictionaries.<sup>15</sup> (Just this reliance makes the occasion-sensitivity in what words with given meanings *say* normally invisible.) That is the default case. In it the ingredients in truth I have identified work in just the way indicated at the start.

Exceptionally, we manifest special intentions as to the understandings our descriptions are to bear, regardless of what they would otherwise be for. Perhaps we sometimes thus get them to bear such understandings. So, e.g., we call the oven hot on some special understanding of its being so, or for some special purpose. Call an oven which is hot on that understanding, or for that purpose, *shmot*. Then when *we* said the oven was hot, we described it as *shmot*. The question now arises: By what standards should the accuracy of a description, on

<sup>15</sup> Typically, we think we know what we would describe correctly. What we thus think we know is *not* whether we would be describing correctly on such-and-such special understanding of being as we described, much less on that understanding conforming to our *beliefs* as to what is describable as thus and so. This last would be no genuine knowledge at all.

that occasion, of an oven as *shmot* be judged? What, on that occasion, would be describable as *shmot*? The story I have told provides the answer.

In any event, however complex exceptional cases may be, default cases are enough to exhibit the ingredients in truth which I meant to highlight here. Such cases are the way we normally get on.

## VII

I have discussed ingredients in truth that deflationism misses. Those ingredients are needed if there is to be any assessing whether the world words speak of is the way their truth demands it must be. What words mean constrains the ways they may speak truth, and the truth that they might speak. It imposes *some* condition on their truth as spoken: the world would have to be, by appropriate standards, as they (in meaning what they do) describe it. The nature of that condition depends on what it would be for a standard to be right. We cannot capture that by speaking of truth in deflationist terms, without referring to occasions for describing, much less saying how what truth is makes them work in fixing standards.

I have not suggested that the work of these ingredients is capturable in a definition; certainly not one which spells out the conditions under which a description would describe truly. Insofar as a deflationist's aim is to reject substantive definitions of truth, that aim may stand, for all I have said.

Deflationists are prone to emphasize that the word 'true' gets much of its importance as a device for 'semantic ascent'. With it, we may characterize propositions as true, or not, without actually expressing them. So, for example, we can quantify, as in 'Everything Jones said is true'. Semantic ascent *is* important. Deflationists have missed, I think, one of the main reasons why it is. To call Pia's words true is to credit her with having described things correctly, as being as they are. That is not to commit oneself to the (present) describability of things by the descriptions Pia used. Pia, say, described a ball as round. But there are many standards for being round. The right ones for her describing need not be the right ones for some describing I might do now. So to say that what Pia said in describing the ball as round is true is not yet to describe the ball as round, or say it to be round. It is not to assert what one would in so describing it.

As for correctness, Pia's words must be judged by the standards appropriate to the circumstances of her speaking them. Those standards need not be the right ones for judging what I would now say, or judge, in describing the ball as round. Part of grasping what it would be for words to be true, or state truth, is grasping what it would be for particular standards to be the right ones by which to judge whether given words did what truth demands what might make given standards

right for judging the truth of some words, wrong for judging the truth of others which give the same descriptions of the same things.

I need to be able to distance myself from another's words in just the way semantic ascent allows, because different occasions call for different standards by which to judge whether things are thus and so. That exposes a side of truth which deflationism misses.

# 4

## Pragmatics

Here are two non-equivalent characterizations of pragmatics. Pragmatics (first version) concerns the linguistic phenomena left untreated by phonology, syntax, and semantics. Pragmatics (second version) is the study of properties of words which depend on their having been spoken, or reacted to, in a certain way, or in certain conditions, or in the way, or conditions, they were. (See e.g. Kalish 1967.)

Here are two equally non-equivalent characterizations of semantics. Semantics (first version) is, by definition, concerned with certain relations between words and the world, and centrally with those on which the truth or falsity of words depends: thus David Lewis's (1972) slogan, 'Semantics with no treatment of truth conditions is not semantics.' Semantics (second version) is defined by this idea: 'A theory of meaning for a language should be able to tell us the meanings of the words and sentences which comprise that language (Platts 1980, p. 2).' So what a semantic theory of English, say, must do is, for each English expression, provide a specification of what it means. Semantics in general would be an account of the nature of such particular theories, or of their subject matter.

Combine these different ideas, and you get a substantial thesis: such things as English sentences *have* storable conditions for truth, and meanings can be given in or by stating these. That *might* be wrong. Perhaps, as J. L. Austin suggested, questions of truth arise at a different level entirely from that of expressions of a language. Perhaps conditions for truth depend, pervasively, on the circumstances in which, or the way in which, words were produced. If so, then on the second version of pragmatics and the first version of semantics, semantic questions are pragmatic ones; whereas semantics (second version), however it is to be done, would have little or nothing to do with truth conditions. Call this the pragmatic view.

This essay argues that the pragmatic view is the right one; that it is intrinsically part of what expressions of (say) English mean that any English (or whatever) sentence may, on one speaking of it or another, have any of indefinitely many different truth conditions, and that any English (or whatever) expression may, meaning what it does, make any of many different contributions to truth conditions of wholes in which it figures as a part. I will first set out the reasons for thinking so, then discuss a few of the most significant consequences.

The issue also emerges in asking what words are for. On one view, bracketing ambiguity, indexicals, and demonstratives, for each declarative English sentence there is a thought which is the one it expresses; its role in English is to express that one. On the pragmatic view this is just what is not so. Independent of ambiguity, indexicality, and so on, what meaning does is to make a sentence a means for expressing *thoughts*—not some *one* thought, but any of myriad different ones. Meaning does that in making a sentence a particular description of how things are, so a means for describing things as that way. Any description admits of many different applications. The same description, applied differently, yields different thoughts. A right application, where there is one, is fixed by circumstances of producing the description, not just by the description itself. If a sentence may thus equally well express any of many thoughts, conditions for the truth of one of these cannot be conditions for the truth of the *sentence*.

### 1. Semantic Properties

There are properties words have, and would have, no matter how we understood them. Being spoken loudly or at 3 p.m. are two. Then there are properties words have, or would have, on one understanding of them, but would lack on another—properties words have, if at all, only in virtue of their being rightly understood in the way they are. I want to consider two classes of such properties.

The first sort of property is one of relating in a given way to truth (or falsity). Properties of being true (false) if, given, of, or only if, thus and so, or thus, or the way things are, are all within this class. (They are all properties words might have on one understanding, and lack on another.) For future convenience, I exclude being true or false *simpliciter* from this class, though I include being true (false) given the way things are. I call these properties *truth-involving*, and any set of them a *truth condition*.

The second sort are properties identified without mention of truth, and on which truth-involving properties depend. Such properties include such things as describing X as Y, calling X Y, saying X to be Y, and speaking of X. The words 'is red', for example, speak of being red and, on a speaking, may have called something red. These properties identify *what* words say. I will call them *content-fixing*, and any set of them a *content*.

One might wonder whether content-fixing properties are not really truth-involving ones in disguise—whether, for example, to call something red is not just to say (of it) what is true of such-and-such things, and true of a thing under such-and-such conditions. In what follows, we will find out whether that is so.

The properties indicated so far might reasonably be called semantic, not worrying overly for the moment about boundaries between syntax and semantics. I will call them that, and any set of them a *semantics*. The latitude allowed here means that not every semantics in the present sense is one words might have.

Some semantic properties may exclude others. Calling something a fish, for example, may exclude, *tout court*, saying what is true of my piano. Call a semantics some words might have *coherent*, keeping in mind that a semantics might thus be coherent on some occasions for speaking, while not on others.

We can raise questions about a semantics, or sort of semantics, without saying which items might have it—whether, for example, English sentences or something else might do so. One thing we may ask of a given semantics is whether it *requires* any further semantics—whether there is a semantics which any words with it must have. Or we may ask whether it is supplementable in a variety of—perhaps mutually exclusive—ways; whether words with it may, for all that, have any of various further semantics.

It is interesting to ask, in particular, whether the semantics an English sentence has in meaning what it does is compatible with any of many supplementations, specifically with any of a variety of truth conditions. To answer that, we need not first say what semantics meaning does confer. We need only find a number of speakings of the sentence on each of which it had whatever semantics its meaning does confer; on each of which, as much as any of the others, those words *did* mean what they do mean. In specific cases we may convince ourselves of that much without knowing just *which* properties meaning confers.

## 2. The Pragmatic View

Is what a sentence means compatible with semantic variety—specifically variety in truth conditions—across its speakings? Consider this sentence:

(1) The leaves are green.

The words ‘are green’, meaning what they do, are means for calling things green. Similarly, meaning what they do, ‘The leaves’, when spoken as in (1), purport to speak of some leaves. What its (present) tense means makes (1), on a speaking, purport (roughly) to speak of things at the time of that speaking. Consider speakings of (1) in which the words did all this, and in all other respects (if any) meant what they *mean*. Does that much semantics require them to have just *one* full semantics on all such speakings? Or is that much compatible with semantic variety, and, specifically, with those words having, on different speakings, any of many truth conditions?

A story. Pia’s Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, ‘That’s better. The leaves are green now.’ She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. ‘The leaves (on my tree) are green,’ Pia says. ‘You can have those.’ But now Pia speaks falsehood.

If the story is right, then there are two distinguishable things to be said in speaking (1) with the stipulated semantics. One is true; one false; so each would



be true under different conditions. That semantics is, then, compatible with semantic variety, and with variety in truth-involving properties. So what the words of (1) mean is compatible with various distinct conditions for its truth.

But is the story right? There are just two grounds for rejecting it. First, one might reject its data by claiming that both speakings of (1), above, share a truth value, require the same for truth, and are true of the same. Second, one might accept the phenomena as presented, but claim that they *are* accounted for by what (1) means—either by some ambiguity in (1), or by some particular way in which what (1) means makes what it says depend systematically on the circumstances of its speaking.

Consider the first option. Either the stipulated semantics makes (1) true of painted leaves, or it makes (1) false of them, *punkt*. If one of these disjuncts is right, appearances to the contrary may be explained in any of a variety of ways. The first task, though, is to choose. Which disjunct is right? One must choose in a principled way. What the words mean must make one or the other disjunct plainly, or at least demonstrably, true.

What we know about what words mean will not solve this problem of choice. Nothing we know about what '(is) green' means speaks to this question: If an object is painted green, should its colour count as what it would be without the paint, or rather as what it has been coloured by painting it? Nor is it plausible that some further development in natural science might resolve this issue. So, it seems, the first option must be rejected. Nor, as we shall see, are colours an unfair example. There are similar problems for any simple predicate, ones left unsolved by what the words in question mean.

We must, then, begin on the second option. Its simplest version is that (1) is ambiguous, or that the words 'are green' are: in one of their senses, they are true of leaves painted green, in another, false of leaves merely painted green. Does 'is green' have such senses in English? I do not think so. But there is a more important question. Suppose it does. Would that yield a different answer to our question about semantic variation?

It would change the answer if the only occasion for saying both true and false things of given leaves in speaking (1) were in case they were painted. But there are indefinitely many more occasions than that provides for saying either of two distinct things in so speaking. Suppose the leaves were not painted (or were painted red), but had a fluorescent green mould growing on them. Or suppose they are painted, but in pointillist style: from a decent distance they look green, but up close they look mottled. Is that a way of painting leaves green? It might sometimes, but only sometimes, so count. So there would be two distinct things to be said in the presumed 'paint counts' sense of 'is green'. And so on.

The above need not be the *only* ambiguity in the English 'is green'. But if words are ambiguous in English, there must be a way of saying just what these ambiguities are; so a fact as to how many ways ambiguous they are. The pair of speakings we considered differed in that each invoked a different understanding

of what it would be for leaves to be green. There is no reason to think that there is any limit to possible understandings of *that*, each of which might be invoked by some words which spoke on that topic. There is not only an understanding on which painting might make it so, but also one on which painting might make it so, as long as it is not in too loose a pointillist style, or too shiny. And so on, *ad infinitum*. If 'green' has, say, thirteen senses, there are, for each of them, various possible (and invokable) understandings of what it would be for leaves to be green in *that* sense. If so, then ambiguity is not a way of avoiding the present conclusion.

It is sometimes said: there is no *uniform* standard for things being green; it is one thing for an apple to be green, another for a tomato to be green, and so on. That idea, though, gets nowhere with the present problem. Throughout the question has been what it is true to say of *leaves*.

Finally, it might be said that the phenomena show 'green' to be a vague term. Perhaps it is in some sense, though we have so far seen no more reason to say so than there is to say the same of any term. But it is hard to see how vagueness is to the point. In one sense, perhaps, words are vague if there is not enough in a correct understanding of them for deciding whether, given the way the things they speak of are, they ought to count as true or false. The *English* sentence (1) is certainly in that condition. But one *speaking* of it may clearly state what is true, while another clearly states what is false. That can only be so if the semantics of (1) on some speakings of it is substantially richer than that fixed for it by the meanings of its constituents, and richer in different ways for different such speakings. So what (1) says on a speaking, of given leaves, etc., is not determined merely by what it, or its parts, mean.

I take the English sentence (1) to illustrate, in the respects noted, what is generally so of a language's sentences—indeed, to illustrate how a sentence of a language *must* function. I have no space for more examples; nor for a satisfying account of why that should be.<sup>1</sup> The reader might anyway test the claim with some further examples of his or her own.

### 3. Domestications

The above, if correct, answers the initial question: what a sentence means, or what its parts do, is compatible with semantic variety; with variety in what such words say or said, and with variety in their truth-involving properties. One might think that compatible with the traditional view, in which semantics is both the study of what words mean and, centrally, of the conditions for their truth; that all said so far is consistent with the meanings of words determining the conditions

<sup>1</sup> For some more discussion see Travis 1989, esp. ch. 1.

for their truth; and even that the general point has long been recognized. One might still think, in other words, that the point may be domesticated within a framework in which what words mean still fixes, in an important sense, what they say wherever spoken. I will discuss two plans for such domestication.

The first plan turns on the idea of ellipsis: some words are to be understood as short for others. A particular 'He'll come', for example, may be rightly construed as a shortened 'He'll come to the party'. Assuming ellipsis were pervasive, how might it help? If (1) may be used to say any of many things, it must, on different speakings, be elliptical for different things: on each it says what that for which it is then elliptical would say. For this explanation to domesticate the phenomena, the things for which (1) is elliptical must not themselves exhibit semantic variation of the sort that (1) did. For example, if a given instance of (1) is elliptical for 'The leaves are green beneath the paint', there must not be more than one thing to be said in *those* words. If the phenomena are as I suggest, this assumption is wrong. I leave this suggestion at that.

The second suggestion revolves around this idea: what words mean *does* determine what they say.<sup>2</sup> But it does not do so *simpliciter*. Rather, it does so as a function of some set of factors, or parameters, in speakings of the words. The parameters allow for different things to be said in different such speakings. Such was always in the plan for linking sentences with truth conditions.

The plan is illustrated by Frege's treatment of the present tense. Frege notes that a speaking of (1) in July might be true, while one in October was false.<sup>3</sup> He observes, correctly, that different things would have been said in each such speaking. One thing this shows is that the tensed verb refers to a specific time or interval, and different ones on different speakings; the words say the leaves to be green at that time.

Frege thought that more was shown. First, that for the present tense the time referred to is always the time of speaking. Second, that where present-tense words are spoken, there is a factor—the time they were spoken—and a function, fixed by what they mean, from values of it to the time they spoke of: in fact, the identity function. So third, that what (1) means determines a function from variables in its speakings to thoughts expressed on those speakings.

Frege's view might be generalized. What *some* words say, or contribute to what is said in using them, varies across speakings of them. Where this is so, the meaning of the words does two things. First, it determines on just what facts about a speaking the semantic contribution of the words so spoken depends. Second, it determines just how their semantics on a speaking depends on these facts. Specifically, it determines a specifiable function from values of those factors to the semantics the words would have, if spoken where those values obtain.

<sup>2</sup> Throughout I leave lexico-syntactic ambiguity aside.

<sup>3</sup> I modify Frege's example slightly. His discussion is in 'The Thought', *Logical Investigations* (Basil Blackwell, Oxford, 1977), p. 27.

The above is a hypothesis. *If* it is true, then while the words (1) may say different things on different speakings, what those words mean determines how they so vary. It determines that the words say thus and so where such-and-such factors take on such-and-such values, for any values those factors may take on (where the thus and so said is what would be true under such-and-such conditions). If that is so, it is reasonable to say that what words mean determines what they say, and when they, or that, would be true. It does so by determining effectively how other facts about their speaking matter to such questions.

But is the hypothesis true? First note that semantics is not history. Sentence (1) will have been spoken only a finite number of times before the heat death of the universe. Suppose that each such time something in particular was said. Then, of course, there is a function from parameters of *those* occasions to what was said in (1) on them. There are many such functions, from many such parameters. That is not semantics. What we wanted to know was: if you spoke (1) on such-and-such occasion (as may or may not actually be done), what *would* you say? The question was whether what (1) means provides an answer to that. The historical remark about actual occasions does nothing towards showing that it does.

The point was that the words 'is green', while speaking of being green, may make any of many semantic contributions to wholes of which they are a part, different contributions yielding different results as to what would count as things being as they are said to be. Are there parameters in speakings of those words which determine just which semantic contribution they would make when? Is there a function such that for each assignment of values to those parameters, there is one particular contribution the words *would* inevitably make, spoken where those values hold? I will not demonstrate here that there are no such things. But there need not be: perhaps for any set of parameters, further possible factors would yield more than one distinguishable thing to be said for fixed values of those.

There are several respects in which the present phenomena are *unlike* central cases where the parameter approach seems promising. One difference is this. In central cases, such as 'I' and 'now', pointing to given parameters seems to be a part of the terms meaning what they do. It is part of the meaning of 'I', and its use in English, that it is a device for a speaker to speak of himself. That suggests speakers as a relevant parameter. If there is no unique semantic contribution, 'I' makes for a fixed value of that parameter, the meaning of 'I' fixes no function from *that* to contributions made in speaking it. By contrast, it is not part of what 'green' means, so far as we can tell, that speakings of it speak of, or refer to, such-and-such parameters. If its contribution, on a speaking, to what is said *is* a function of some parameters—say, implausibly, the speaker's intentions (see Ch. 2 above)—saying so is not part of saying what 'green' means. The parameter approach does not *automatically* suggest itself here as it did with 'I'.

This difference between 'I' and 'green' shows up when it comes to saying what was said. Consider a speaking of the words 'I am in Paris'. Ignore any possibilities for various contributions by 'in Paris', or by the present tense at a time. Then, knowing nothing more about the speaking, we know that, in it, it was said *that* the speaker, whoever *s/he* may be, was, at the time of speaking, whenever that was, in Paris. However in the dark we may be on those points, we *do* thus specify which fact (or non-fact) was stated. Not so for speakings of (1). Suppose that Pia spoke those words, and that *we* say of that, 'Pia said that the leaves she spoke of were, at the time of speaking, green.' We will not have said *what* Pia stated unless *our* 'green' made some definite contribution to what *we* said about Pia. But, as we have seen, 'green' may make any of many contributions of the needed sort. If it made one such in *our* words and a different one in Pia's, then what we said about her is *false*. We may, for example, have said her to say what would be false of green-painted leaves, while what she said would be true of that. The information contained in the meanings of the words she used is thus not enough for specifying, however uninformatively, *which* fact (or non-fact) she stated.

In speaking (1) literally, one does what then counts as calling leaves green. That may be one thing that sometimes counts as 'saying that the relevant leaves were green'. But *such* a use of 'say that', if there is one, does not purport to specify which fact (or non-fact) was stated. It says nothing that allows us to associate what was said with a truth condition for it. So it does not point to a function, fixed by meaning, from speakings to thoughts expressed in them.

A second contrast between present phenomena and such things as 'I' and 'now', traditionally conceived, is suggested by this remark of Frege's:

the content of a sentence often goes beyond the thought expressed by it. But the opposite often happens too; the mere wording, which can be made permanent by writing or the gramophone, does not suffice for the expression of the thought. . . . If a time indication is conveyed by the present tense, one must know when the sentence was uttered in order to grasp the thought correctly. Therefore the time of utterance is part of the expression of the thought. . . . The case is the same with words like 'here' and 'there'. In all such cases the mere wording, as it can be preserved in writing, is not the complete expression of the thought; the knowledge of certain conditions accompanying the utterance, which are used as means of expressing the thought, is needed for us to grasp the thought correctly. Pointing the finger, hand gestures, glances may belong here too. (Frege 1918–19, trans. pp. 10–11)

We begin with the idea that sentences are related to thoughts in this way: for each sentence there is a thought which is the thought it expresses.<sup>4</sup> With indexicality, we lose that idea. There is no particular thought which is the one the sentence 'I am here' expresses. Perhaps, though, we may regain that idea if we permit ourselves to generalize the ordinary notion of a sentence. Ordinarily, we think

<sup>4</sup> Once again, ignore lexico-syntactic ambiguity.

of a sentence as a string of words. Suppose, though, we drop that idea. Let us call something a symbol if it has two features. First, it is individuated by purely non-semantic features, as a word might be individuated by its shape.<sup>5</sup> Second, it has semantic properties, where we will take that to be so if it makes a definite, specifiable semantic contribution to the whole, or wholes, of which it is a part. We might regard a (generalized) sentence as a structured set of symbols in this sense. So, if Frege is right about its semantic contribution, a time of utterance may be a symbol, and hence a constituent of a sentence in this sense. An utterance ‘The leaves are green’ in July would then count as a different sentence from an utterance, ‘The leaves are green’ in October—an odd, but coherent way to speak.

If the only deviations from the rule that, for each sentence, there is the thought it expresses are represented by the sort of case Frege has in mind, then we may now regain the initial idea in this form: for each *generalized* sentence, there is a thought which is the thought it expresses. But the phenomena exhibited by (1) cannot be domesticated in this way. There is no identifiable feature of a speaking of (1) which counts as a symbol in the present sense, and whose semantic contribution to the speaking is identifiable with precisely the set of truth-involving properties (1) would have so spoken. If the phenomena (1) exhibits are pervasive, then even a generalized sentence, no matter what extra symbols it contained, might be used to say any of many things.

Wittgenstein held that any symbol is open to different interpretations; and that under different circumstances, different identifications of its content would be correct. That is the moral of his discussion of rules and what they instruct (*Investigations*, §§84–7). His arguments apply as well to generalized symbols as to others. If he is right, then the demonstration omitted here, that the parameter approach *cannot* work, is anyway to be found.

#### 4. Implicature

Suppose that I were the doctor and a patient came to me, showed me his hand and said: ‘This thing that looks like a hand isn’t just a superb imitation—it really is a hand’ and went on to talk about his injury—should I really take this as a piece of information, even though a superfluous one? (Wittgenstein 1969, §461)

I am sitting with a philosopher in the garden; he says again and again, ‘I know that that’s a tree,’ pointing to a tree that is near us. Someone else arrives and hears this, and I tell him: ‘This fellow isn’t insane. We are only doing philosophy.’ (Ibid. §467)

Wittgenstein cites some bizarre things to say. We do not say such things, barring very special occasion to do so. But what does that mean? Suppose one says them anyway. Despite the oddity, might one have spoken truth?

<sup>5</sup> Strictly speaking, this is false of words (consider e.g., homonyms). But ignore that for now.

The philosopher does acrobatics recklessly close to the tree. ‘That’s a tree over there,’ someone warns. ‘I know that’s a tree,’ he replies testily. ‘Well, then, shouldn’t you be more careful?’ Here the philosopher speaks truth. So, one might reason, he *does* know these things. But one cannot cease to know things, or so it seems, *just* by moving from one conversation to another. So however bizarre saying so may be in other cases, for all that, he speaks truth there too. So one might reason.

But this is a bad argument. For it *may* be that words like ‘I know I’m wearing shoes’ vary their semantics from speaking to speaking. If some speakings of them speak truth, that does not mean that all will. We cannot generally reason: Pia spoke truth when she called the leaves green; so if I call them green, I will speak truth too. That was the moral of section 2. The point would be, not that the philosopher ceases to know something by changing conversations, but rather, that on one occasion he counts as knowing such-and-such, on another not.

There is, though, a form of account on which many bizarre things we ‘would not say’, would, for all that, be true. The idea is due to H. P. Grice. The starting point is the observation that saying is only one of numerous ways for words, or speakers of them, to represent things as so. There is also implying, suggesting, insinuating, presupposing, and so on. *That* insight did not originate with Grice. Grice, though, concerned himself with a particular class of such representations, which he called implicatures, using the verb ‘implicate’ for the sort of representing in question. Implicatures come in two sorts: conventional and conversational. Conventional implicatures are features of the meanings of the terms involved. They are illustrated by ‘Pia dissuaded Tod from leaving’, and ‘Sam struggled to reach the lectern’. The first represents Tod as at least having thought of leaving; the second represents Sam as facing some obstacle to reaching the lectern. But the first does not *say* that Tod had thought of leaving, nor the second that there was an obstacle. That does not yet mean, that, for example, the second might be *true* were there no obstacle. It leaves it obscure what could make it so. But it may facilitate arguing the point. In any event, just as to use ‘It’s green’ to mean what it does *is* to call something green, so to use ‘struggle’ to mean what it does, in a case like the above, is to suggest or imply that there is an obstacle. Grice suggests that it is difficult to produce words with a conventional implicature without implicating that. Such implicatures are not, or hardly, what Grice calls ‘cancellable’. That he takes to be a main identifying feature of them.

Some implicatures, Grice notes, arise only on certain speakings of words, so *are* cancellable. These Grice calls conversational implicatures, and he explains them thus (though in much greater detail than given here). In normal conversation, we represent ourselves as observing certain maxims, and may be supposed to do so. Grice calls these *conversational maxims*. Examples are: be co-operative, be brief, be informative, and be relevant. Sometimes a speaker *seems* to violate some of these maxims. But it may be that he would not have if such-and-such, and it may

be unreasonable to take the speaker to be violating them. We may then reason thus. The speaker said that P (in saying 'W'). Saying P (or saying it in 'W') would violate the maxims unless Q. The speaker was not violating the maxims. So (according to him) Q.

A speaker may intend for us to avail ourselves of some inference of this sort, to a given conclusion that (according to him) Q. It may be part of the proper understanding of his words that he so intends. In that case, the speaker has, or his words have, conversationally implicated that Q. For example, Pia may say, 'Jones submitted a sequence of English sentences, divided into paragraphs, and titled "What is truth?."' If this is merely a way of saying that Jones submitted an essay, then it violates the maxim of brevity. Pia would not do *that*. So, by the suggested sort of inference, we may conclude that there is, according to Pia, something which distinguishes Jones's work from a proper essay—perhaps its incoherence. It may have been given to be understood that we were so to reason. In that case, the conclusion was conversationally implicated.

The notion of conversational implicature points to a particular sort of understanding some words, on some speakings, may bear. Nothing in the pragmatic view suggests that there should not be such understandings. Note, though, that, as Grice insists, for Q to be conversationally implicated in words 'W', Q must follow from what 'W' said, or the fact that 'W' said it, or both. So we might ask what Grice thinks words say. He is quite clear about that:

In the sense in which I am using the word *say*, I intend what someone has said to be closely related to the conventional meaning of the words (the sentence) he has uttered. Suppose someone to have uttered the sentence *He is in the grip of a vice*. . . . One would know that he had said, about some particular male person or animal x, that at the time of the utterance . . . either (1) x was unable to rid himself of a certain kind of bad character trait or (2) some part of x's person was caught in a certain kind of tool or instrument . . . But for a full identification of what the speaker had 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* [a decision between (1) and (2)]. (Grice 1989/1975, p. 25)

This is just the rejected conception of saying. On it, for example, bracketing lexico-syntactic ambiguity, we can always form a guaranteed-true report, in indirect speech, of what was said in any arbitrary speaking of given words: if the words were 'The leaves are green', then that the relevant leaves were, at the relevant time, green. To think that is to miss the possibility of occasion-sensitivity in the content of 'green'. So Grice's conception of saying cannot be assumed in any argument directed against an instance of the pragmatic view.

Grice aimed to resuscitate views fallen into disrepute, largely through what were, in effect, early applications of the pragmatic view. For example, the idea of conversational implicature was first developed specifically in aid of reviving some notion of a sense datum. With that in mind, let us return to the bizarre remarks



with which this section began. Consider 'I know that that's a tree'. It would usually be bizarre to say that, for example, where the tree was in plain view and no doubt of any kind had arisen as to whether it was a tree. Grice invites us to entertain the possibility that the reason we would not say such a thing in such circumstances is that if we did, we would conversationally implicate something not so. He means that idea to encourage us to ask whether what would be said if one did so speak is anyway something true, or rather something false; and to expect one choice or the other to be *correct*.

In using 'know' bizarrely we *may* conversationally implicate something (though there is a problem if conversationally implicating that Q absolutely requires saying that P). But the pragmatic view offers another explanation of why, in some situations, we would not say 'I know that . . .'. Suppose that 'know' may make any of many distinct semantic contributions to wholes of which it is a part, and varies its contribution from one speaking to another. Then, describing someone as he is at a time, we would, on some occasions, say something true in saying him to know that X is a tree, and, on other occasions, say something false in saying *that*. For there are various things to be said in so describing him. In that case, circumstances of a speaking of 'N knows . . .' may confer on it a supplement to the content provided by the meanings of the terms alone. For some such supplements, the result will be stating truth; for others it will be stating falsehood. But *some* circumstances may fail to confer a supplement of either of these sorts. Words produced in such circumstances would have a content still supplementable in either way. But a content still so supplementable can require neither truth nor falsity. Speak, in those circumstances of N knowing that it's a tree, and one will fail both at saying what is true and at saying what is false. Nothing either so or not-so will have been said to be so. Recognizing that, where it is so, may make one refrain from so speaking. In that case, the idea, encouraged by Grice, that if we said it anyway we would at least say something true or else something false, is simply a mistake. In that case, conversational implicature could not be a consequence of the fact of having said *that* such-and-such. There is no such fact.

That the content of words is consistently supplementable in more than one way is not in itself a block to those words stating truth. It is so only where different such supplements, or different ones within some range of reasonable ones, yield different results as to truth—where, that is, the content to be supplemented is compatible both with truth and with falsity. So it just *might* be that if you say irrelevantly, pointing at your brogues, 'Those things are shoes', there is no compelling reason to deny that you have spoken truth (though the situation changes if you are wearing four-eyelet low moccasin boots, or even just moccasins). That is typically not how it is for philosophically sensitive terms like 'know'. That is one lesson the long history of scepticism teaches us. (If there must be an *occasion-insensitive* answer, just when *does* someone count as knowing there is a tree before him?)

This last point shows the problem in applying the notion of implicature where it is meant to carry philosophic baggage, notably where it is meant as a way of dismissing claims about what ‘we would not say’ as philosophically irrelevant. Where those claims point to occasion-sensitivity, they are philosophically highly relevant. It is all very well to insist, for example, that either Sam does or doesn’t now know that he is wearing shoes, full stop; and that if you said, bizarrely, ‘Sam knows he is’, you would either state truth or state falsity. Sooner or later, though, one must choose. Which is it? If, applying the pragmatic view, we carefully assemble a perspicuous view of the *different* things we at least take ourselves to say to be so, on different occasions for speaking of Sam, in saying him to know precisely that, then *either* there is a principled way of choosing between them (or choosing a further candidate) by appealing to what is recognizably so about what ‘know’ means, *or* they show that no one answer to the question is the right one occasion-independently. Prospects for the first alternative are dim.

## 5. Metaphysics

The English ‘is green’ speaks of a certain way for things to be: green. One might say that it speaks of a certain property: (being) green. If we do say that, we must also say this about that property: what sometimes counts as a thing’s having it sometimes does not, so that there are, or may be, things which, on some occasions for judging, count as having the property, and on others do not. If for a property to have an extension (at a time) is for there to be a definite set of things (at that time) which are just those things (then) with that property, then this property does not have an extension, even at a time. Better put, it makes no sense to speak of ‘its extension’.

Is all this just vagaries of the English ‘is green’? Two related questions arise. First, might there be predicates which did not vary their contributions to what was said with them in the way that ‘is green’ does? If we said such a predicate to speak of a property, that property *would* have an extension, at least at a time. Such a predicate could not vary its contributions to wholes so that, in ascribing that property to an object (at a time) it would be possible to speak truth and also possible to speak falsehood. So there would be no call for saying of anything that it sometimes counted, and sometimes didn’t, as having (at a given time) that property. Second, can we preserve the idea that (genuine) properties have extensions by supposing that predicates like ‘is green’ simply refer to different properties on different occasions (and that it is by their thus varying their referent that they make different contributions to different wholes)?

Why might one want properties to have extensions? First, one might think that we can gain this for properties by definition—by ‘property’ we just mean what has an extension—and that extensions are convenient means for counting properties (as one or two). Second, one might take such a view of properties as

mere sane realism. We cannot change, say, the way a cow is by thinking about it. As a rule, the cow stays just as it is no matter how we think of it. We may read, or misread, that sane thought thus: those ways for things to be which are, or count as, ways the cow is count as ways the cow is no matter how we think about the cow, or them. So for any genuine way for things to be, either the cow is that way (at a time), or it is not, *punkt*. The same goes for any other object. In which case genuine ways for things to be have extensions (at times). But whatever there is in favour of this line of thought, I suggest that both our questions merit negative answers.

I begin with the first. I will state the main point, though there is here no space for detailed argument. Once we fix what 'is green' speaks of—green—we then note that there are different possible understandings of what it would be for an object (or some objects) to be *that* way (green). These are possible understandings in that they represent what one *might* regard as a thing's being green. So, for each, some item may be said, in calling it green, to be green on *that* understanding of its being so. And for each, that may be the *right* understanding (on some occasion) of what being green would come to. 'Is green' provides a particular description for things, expresses a certain concept. What is said in using it depends not only on what that description is, but on how that description, or that concept, is, or would be, applied in fitting it to particular circumstances of its use.

Suppose, now, that we identify an understanding of being green—say, the understanding on which an item was said to be green in some particular speaking of 'is green'. We now introduce a predicate—say, 'is green\*'—which, by stipulation, is to mean *is green on that understanding of being green*. This predicate speaks, as it were, of a finer-grained property than 'is green' (as such) does. May *this* predicate make different contributions to what is said in wholes of which it is part? It may if there are different possible understandings of what it would be to be green on that understanding; two different things to be said as to whether such-and-such *is* being green on that understanding of what it would be to be so. As far as we can tell, this always will be so. We understand, for example, that paint is to count as changing colour, and not as hiding it. We now encounter a rather poor paint job: you *could* say that it covered the original colour, but you could view the original colour as still showing through enough that the object had not yet been made the colour of the paint, even on the indicated understanding of its being that colour. An understanding of being green, insofar as we can identify one, seems unable to foreclose in principle on the possibility of differing, but, apart from particular surroundings, equally sane and sensible views of what *that* understanding entails.

A predicate about which the pragmatic view was wrong would be one which did not admit of different possible understandings of what it would be for some item to fit the description which that predicate provides (or for the description to fit some item). The right understanding of it would foresee every eventuality in or to which the description might be applied. There is reason to think that

no such predicate is available to human beings, at least given the way we in fact cognitively conduct our affairs. Again, what is said in applying a given description depends on *how* it is applied, and how, in given circumstances, it ought to be.

Now for the second question. First, if the first point is correct, then no understanding *we* could have of being green, so none that might attach to a particular use of 'is green', would be one on which 'is green' spoke of a property, if a property must have an extension. To paraphrase Wittgenstein, we refine our concepts, or understandings, for particular purposes—so that *in fact*, in the situations we face or expect, unclarity as to what to do or say does not arise. In doing that we neither reach, nor aim at, that absolute clarity on which we would speak of what had definite extensions. Where 'is green' has made different contributions to different wholes, we may identify different things for it to have spoken of each time—being green on this understanding, and being green on that one. So we may see the predicate as varying its reference across speakings of it. But we must not mistake these different things for properties with extensions. Second, if we cannot have a predicate for which the pragmatic view does not hold, then, equally, we have no means for specifying properties to which extensions may sensibly be ascribed. In any event, the phenomenon we have to deal with is not merely that predicates vary their contributions to wholes, but also that, whatever a predicate may be said to speak of—being such-and-such—what would sometimes count as an item's being *that* other times would not.

## 6. Perspective

Given words may have any of many semantics, compatibly with what they mean. Words in fact vary their semantics from one speaking of them to another. In that case, their semantics on a given speaking cannot be fixed simply by what they mean. The circumstances of that speaking, the way it was done, must contribute substantially to that fixing. As pointed out earlier, this does not mean that there is a function from certain parameters of speakings to semantics, taking as value for each argument the semantics words would have where those values held. It thus also does not mean that there might be a precise theory, generating, for each semantics words might have, necessary and sufficient conditions for their having that. Still, we may describe how circumstances do their work.

Here is one thought. The words 'is green' are a means which English provides for calling things green (describing them as green, etc.). If, in speaking English, you want to call an item green, those words will do. Speak them literally, seriously, and so forth, and you will then count as having done just that. The truth of what you say in calling an item green should turn precisely on whether the way that item is then counts as its being green. These two remarks jointly identify which truth-involving properties any such words must have: they are

true of, and only of, those ways for things to be which counted, at their speaking, as the item they spoke of being green. Similarly for other English predicates.

Where you called an item green, the truth of your remark turns on whether it *then* counted as being green. On different occasions, different ways for an item to be would count as its being green. That variation means that, on different occasions, calling an item green will confer different truth-involving properties on your words. Consider two occasions which differ in this respect. On each, words which call an item green will have some set of truth-involving properties, which is, therefore, a possible set of such properties for words with that content to have. Those truth-involving properties, and the property of calling that item green, cohere on at least some occasions for so describing things. But those truth-involving properties *cannot* be those of words with that content produced on the other. That would not correspond to what, on the other, counts as something's being green. So each of the above semantics, available as it is on some occasions, is unavailable on others. I can sometimes speak truth in calling painted leaves green; but I cannot do so in circumstances where their being so painted does not count as their being green.

Let us pursue this thought. Consider:

(2) Today is a sunny day.

Spoken on day D, (2) would, typically, speak of day D. It would also identify the day it speaks of in a particular way: it speaks of that day as the day of its speaking, and represents it as identified by that fact. Since some speaking of (2) has both the semantic properties just mentioned, the two jointly form a semantics which is at least sometimes coherent. Let  $D^*$  be the day after D. Words produced on  $D^*$  could not have the semantics just mentioned. They could not speak of D and say it to be sunny while, on their proper understanding, identifying the day they speak of as the day of their speaking. On day D, we may express, or think, a thought with both those features. On other days (in normal circumstances) we cannot. Let us say that words with a semantics which is only sometimes available, in the above sense, express a perspectival thought, and have a perspectival content.

Now the point of the discussion of 'is green' may be put this way. Perspectival thought is the normal and pervasive case. On one occasion, we call an item green (at a time), and thereby produce words with such-and-such truth-involving properties. On another occasion, we may, if we like, say the same item to be green (at that same time). But our doing that may require that our words have quite different truth-involving properties. Those of our first remark may not correspond to what would count, on the occasion of this further speaking, as that item's being green. If that is right, it is fair to suppose that perspectival thoughts are the typical sort of thoughts we think. One might say: we relate cognitively to the world in essentially perspectival ways.

Now consider two minor puzzles. First, I have said there is something true, and also something false, to be said of given leaves, and their condition at a given

time, in saying them to be green. How can this be? Consider the true thing to be said. What could make it true, other than the fact that the leaves are green? But, if that is a fact, how could one speak falsehood in saying no more nor less than that about them? Second, if there *are* those two things to be said, then *say* them, or rather, state the true one and deny the false one. To do so, you would have to call the leaves green, and then deny that they are that, as in ‘The leaves are green, and the leaves are not green’. But that is a contradiction, so cannot be true. So what the pragmatic view requires that it be true to say is something it could not be true to say. So the view is wrong.

The first puzzle’s rhetorical question has a non-rhetorical answer. What could make given words ‘The leaves are green’ true, other than the presumed ‘fact that the leaves are green’, is the fact that the leaves *counted* as green on the occasion of that speaking. Since what sometimes counts as green may sometimes not, there may still be something to make other words ‘The leaves are green’ false: namely, that on the occasion of *their* speaking, those leaves (at that time) did not count as green.

As for the second puzzle, we are challenged to say something literally unsayable—*not*: sayable-but-false, but rather not sayable at all. We ought to decline the challenge. On some occasion, words which call given leaves (at a time) green may (thereby) have truth-involving properties in virtue of which they are true. On some other occasion, words which deny those same leaves to be green may similarly be true. But given the way (described above) in which occasions work to forge a link between content-fixing properties and truth-involving ones, there is no occasion on which both these feats could be accomplished at once; so none on which ‘The leaves are green and the leaves are not green’ could have the semantics which a conjunction of those two truths would have to have. If the occasion is one on which the way those leaves are counts as their being green, then no words could have the semantics of the true denial; and *mutatis mutandis* if on the occasion the way the leaves are does not count as their being green. Each of the thoughts provided for above is a *perspectival* thought; and, in virtue of its perspectival character, unavailable to be expressed at all on any occasion on which the other is expressible.<sup>6</sup> The nature of semantic variation thus allows us to decline the challenge.

These are banal examples. In philosophy, neglect of perspectival thought often leads to more excitement. A philosopher may sense, for example, that our concepts apply as they do against a background of our natural reactions; if we naturally viewed things *quite* differently, we might apply the concepts we now have so as to speak truth in saying what it would not now be true to say. Asked to express some such truths, the philosopher is reduced to nonsense. Naturally

<sup>6</sup> More precisely, any occasion on which a thought with the semantics of the first is expressible is *ipso facto* one on which a thought with the semantics of the second is not. I do not mean to prejudge questions of thought-identity.

enough. He was describing other perspectives. Some things said truly from them are not so much as expressible at all from his own.

## 7. Thoughts

Frege writes,

Without offering this as a definition, I mean by 'a thought' something for which the question of truth can arise at all. (1918–19, trans. p. 4)

Thoughts, for Frege, are not words. For him words are true only in a derivative sense: just in case they express a thought which is. For words are always open to, and in need of, interpretation. They are true, if at all, only on a given understanding of them (even if it is their proper understanding). Words 'Mary had a little lamb' may be a remark on husbandry, or one on gastronomy and, perhaps, true if understood the first way, false if understood in the second. Truth and falsity seem to correspond to understandings words may have, rather than to the words themselves (which Frege conceives as a quite different matter). It is the understandings, as opposed to the words, which settle questions of truth and falsity. So, on his view, it is for understandings, and not for words, that questions of truth and falsity arise. Words, apart from an understanding, could not be true or false at all.

If words admit of interpretations, then conceivably they may bear different understandings on different occasions for understanding them. Such shifts in interpretation could bring with them shifts in truth value. So if words were the primary objects for which questions of truth arose, it would be conceivable, for any sort of semantic object, that one and the same item should count as true on one occasion for assessing it, false on another.

Thoughts, for which questions of truth are, strictly speaking, to arise, are meant to be free in principle of both of the above features. They are to be absolutely immune to interpretation; and they are to be true or false absolutely, independent of the ways, if any, in which they enter into our thinking. On Frege's view, only such semantic objects could be material for logic.

We may extend the notion of semantic property so that thoughts have a semantics too. The semantic features of a thought will be just those features by which one thought may be distinguished from another. Among these will be such things as being about eating ovine, and such things as being true if Mary ate a bit of ovine, hence, on the above plan, both truth-involving and content-fixing properties. Its truth-involving properties are meant to be just those its content requires. Moreover, it is meant to have all this semantics intrinsically: any thought, no matter how encountered, is that thought iff it has that semantics. This means that the content of a thought—unlike the content of words—must determine its truth-involving properties *inexorably* (to coin a term): there are

no two sets of truth-involving properties such that an item with that content might have the one but not the other, and also vice versa; there is *one* set of truth-involving properties which is *the* set any item with that content *must* have. For if not, then a thought's having that content might, on some occasions, make it count as having one set of truth-involving properties, and on others make it count as having another, counter to the tenet that every thought has its truth-involving properties intrinsically.

Why must thoughts have inexorable content? Suppose C is a non-inexorable content. Then there might be an item with C and truth-involving properties T, and an item with C and distinct truth-involving properties T\*. But truth-involving properties are meant to be those which content requires. So these must be two items differing in further content-fixing features. This means that an item with C is, so far, open to interpretation: it might, for all that, bear any of several distinct understandings. That is to say: it might, for all that, be, or (if words) express, or represent, any of several distinct thoughts. So C is not the (whole) content of a thought.

Thoughts are identified precisely by their semantics, whereas words are identified by shape, syntax or spelling, or by the event of their production. The identity of words leaves their content open. So the content of given words must depend on further factors: on the character of their surroundings. This leaves it open that their surroundings might, on some occasions of considering them, count as conferring one semantics on the words, while on other such occasions those surroundings might count as conferring another. In that way, the semantics of words—how they are rightly understood—may be an occasion-sensitive affair. By contrast, the semantics of a given thought is meant to depend on *nothing*. So there are no such possibilities for variation across occasions in the semantics a given thought *counts* as having.

Thoughts, as thus conceived, are not open to interpretation. They are what Wittgenstein called 'shadows': semantic items interpolated between words and the states of affairs that make words true or false, and somehow more closely tied to those states of affairs than mere words could be. About shadows, Wittgenstein said:

Even if there were such a shadow it would not bring us any nearer the fact, since it would be susceptible of different interpretations just as the expression is.<sup>7</sup>

How could this be true of thoughts? Could thoughts admit of interpretation? If so, how?

There are too many strands in our inherited notion of a thought to unravel them here. But here is a sketch of a framework for relevant issues. To begin, one *might* think to buy the semantic absoluteness of a thought—its immunity to interpretation—by stipulation. Wherever I would say something to be so in

<sup>7</sup> Reported by Moore 1993 (1954), p. 59.



saying 'S', and it is determinate what, I may, it seems, refer to a thought in saying 'the thought that S'. I may also say, correctly, it seems: 'The thought that S is true iff S'. In saying that, I ascribe a set of truth-involving properties to the thought I refer to; in fact, whatever such properties my words 'S' then had. For I say the thought to be true exactly where what is so according to my words 'S' is so. So, it seems, we might stipulate that the thought I thus refer to is precisely the one with those truth-involving properties.

This is not quite enough. A thought cannot *just* have truth-involving properties. It must have a content. What content should that be? Here we come up against another strand in the conception of a thought. A thought is meant to be something that can be expressed in various words, or speakings, on various occasions. If you now express a thought, I can later express that very thought virtually whenever I like. On any plausible version of that view, words W and W\* may express the same thought while differing in content. Frege gives this example:

If someone wants to say today what he expressed yesterday using the word 'today', he will replace this word with 'yesterday'. Although the thought is the same, its verbal expression must be different in order that the change of sense which would otherwise be effected by the differing times of utterance may be cancelled out. (1918–19, trans. p. 10)

The word 'today' brings with it a different contribution to content than the word 'yesterday'. Frege's two sentences are not alike in content-fixing properties. Yet, for good reason, Frege takes it that the one sentence, produced under certain circumstances, would express the same thought as the other sentence produced under certain others. If so, then the content-fixing properties of that thought are liable to vary across occasions.

The question is: just *how* may content vary while words express the same thought? One idea would be that W and W\* express the same thought only if they apply the same concepts to the same objects. But this will not do. It does not even allow for Frege's example. It collapses completely if we return to the notion of perspective. On some occasions, in calling given leaves green one would state truth; on others, in calling those leaves green one would state falsehood (and not because the leaves changed). Apply a given concept to the leaves in different surroundings, and you will produce words with very different truth-involving properties. The semantics of some such words, produced in given surroundings, is unavailable in other surroundings for any words. Words with the *content* of those words, in the other surroundings, may have truth-involving properties so different that, at least for some purposes, we cannot take them to have expressed the same thought. The false remark about the leaves, for example, was not the same thought as the true remark. So if, in the changed surroundings, one wants to express the same thought again, one must *not* speak of the same concepts and objects. What it would take to express the same thought again is nothing more nor less than an adequate paraphrase. If the original words were 'The leaves are

green', then, depending on surroundings, an adequate paraphrase might be 'The leaves are painted green'.

There is no space here for an account of what makes paraphrases adequate. But here are two remarks. First, adequate paraphrases may need to share crucial or relevant truth-involving properties; but they are unlikely to share *all* truth-involving properties. In remote enough circumstances, leaves may be green in the sense in which they were said to be in a given 'The leaves are green', but not painted green (perhaps dyed); though, for current purposes, 'The leaves are painted green' was an adequate paraphrase. Second, suppose on an occasion I express a thought in saying 'The leaves are green'. Then whether, on another occasion, words W are an adequate paraphrase of what I said may well depend on the occasion for the paraphrase, and perhaps, too, on the occasion for considering that occasion.

Thoughts viewed from this position lose their claims to have some *one* semantics intrinsically, and to be immune to interpretation. If, with perspective in mind, we ask what would count as producing some given thought again, and if we consider all the occasions for posing that question, we see how *that* thought may count on some occasions as having semantics which it would not count as having on others. For it may on some occasions admit of paraphrases it does not admit of on others. Nor need it ever have an inexorable content. To see how thoughts admit of interpretation, one need only know how to look for occasions for interpreting them.

## 8. Concluding Remarks

There is much left to discuss, but no space left to discuss it. It is thus time to commend the subject to the reader. The pragmatic view gives a substantially different form to virtually every philosophic problem, not just in philosophy of language, but wherever puzzles arise. The new form may make some of these problems more tractable. For a start we will need new conceptions of logical form, and of such related notions as intensionality. These may yield new things to say on such questions as whether 'if-then' is transitive. We may then take a fresh look at what we say of people in ascribing propositional attitudes to them, and at understanding itself. Such a look, I predict, would make philosophy of psychology take a fresh course. It is also worth a look, from the pragmatic view, at problems of knowledge, of explanation, of freedom and responsibility, and so on. Some of this work is begun. There is much left to explore.

# 5

## Sublunary Intuitionism

Truth matters to all those capable of thinking falsely (so of *thinking* things); but to philosophers in a special way. Philosophers have spoken volumes on the topic, only occasionally memorably. One such occasion is Michael Dummett's early essay, 'Truth'. That essay sets out, with economy and clarity, a position which is original, radical, and deep. Dummett has since done much to develop that position. But the present essay concentrates on that one elegant statement of it.

In his essay Dummett defends two intuitions. Each has its exponents. Few, though, defend both. If both are right, that *may* argue for a view I dub *sublunary intuitionism*: in an important sense (though one yet to be clarified) the logic of ordinary discourse is intuitionist, not classical.

I will argue, though, that Dummett's first intuition is incorrect. His second intuition stands, on a more radical, perhaps more simple-minded, reading than he gives it. But it is the tension between the two intuitions, if anything, that argues for sublunary intuitionism. Neither alone makes the case. Or so I argue.

### 1. What Truth Is Like

Dummett's first intuition, roughly, is that no statement could be neither true nor false; or it could never be right to say so. The second is that it is not so of every statement that it is guaranteed to be either true or false. The core idea behind the first intuition is, as Dummett later put it:

The question is whether there is a place for a convention that determines, just by the meaning of an assertoric utterance of a certain form, that, when all the relevant information is known, the speaker must be said neither to have been right nor to have been wrong: and it seems clear that there is no such place. (1978, p. xviii)

In order to fix the sense of an utterance, we do not need to make two separate decisions—when to say that a true statement has been made and when to say that a false statement has been made. (1959, p. 10)

From which Dummett takes it to follow that:

A statement, so long as it is not ambiguous or vague, divides all possible states of affairs into just *two* classes. For a given state of affairs, either the statement is used in such a way

that a man who asserted it but envisaged that state of affairs as a possibility would be held to have spoken misleadingly, or the assertion of the statement would not be taken as expressing the speaker's exclusion of that possibility. If a state of affairs of the first kind obtains, the statement is false; if all actual states of affairs are of the second kind, it is true. It is thus *prima facie* senseless to say of any statement that in such-and-such a state of affairs it would be neither true nor false. (1959, p. 8)

A statement represents itself as saying how things are; not as *perhaps* saying how things are if conditions are thus and so, but otherwise going in for some other line of work. Unless the way it represents things is the way things are, it is no success in the terms it set itself. If things are a way that fails to count as how it represented them, then *it* has failed at just the task of representing rightly that it set. Such failure, Dummett suggests, is falsity. So no way things might be could make it so that the statement was both not true and not false. If a statement could be neither true nor false, Dummett suggests, that would have to be in virtue of some special convention—some understanding as to what it was to do—setting out the circumstances, and the way, in which it might do that. But, Dummett says, and I agree, there is no such convention; a statement bears no such understanding.

Many philosophers find themselves in Dummett's neighbourhood on this first point. Timothy Williamson is one. Williamson realizes that not everything is either true or false—my hat, for example. What is true or false must at least represent things as being thus and so. But what does that much, he thinks, cannot fail to do it either truly or falsely. So Williamson proposes this principle:

If *u* says that *P*, then either *u* is true or *u* is false. (1994, p. 187)

This principle, he argues, brooks no counterexamples; for a counterexample would entail a contradiction.<sup>1</sup> So, for example,

[I]t is incoherent to suppose that vague utterances in borderline cases both say something and fail to be either true or false. It is coherent to suppose them to be neither true nor false only at the cost of treating them as though they said nothing. (1994, p. 198)

Williamson emphasizes that

The argument does not immediately show that bivalence must be asserted for particular utterances, only that it must not be denied. (1994, p. 193)

But Williamson has no interest in refraining from asserting that. Indeed, his own view of vague statements depends on it. That means that Williamson flatly does not share Dummett's second intuition. So any problem there may be in squaring Dummett's second intuition with his first is not one Williamson faces.

<sup>1</sup> Williamson's argument is serious. My position needs urgently to disarm it. But editorial considerations dictate that I do not respond here. I hope to do so in another place.

Now for Dummett's second intuition. This begins with the thought that

A statement is true only if there is something in the world *in virtue of which* it is true. (1959, p. 14)

Not every statement, Dummett holds, is *guaranteed* to fulfil that condition. Consider a statement, of Jones (deceased), that Jones was brave. Jones led a quiet life, without occasion for either bravery or cowardice. One might have various theories of bravery. Perhaps there is a bravery gene, or some other mechanism—neurology, say—that distinguishes unevincing bravery from the lack of it. And perhaps Jones had such a mechanism or lacked it. But then again, perhaps not. For if there is *any* such mechanism, that is only contingently so. Absent such a factor, there might, we may suppose, be nothing to make it right to call the statement true, and nothing to make it right to call it false. Or at least nothing would make it right for us to suppose otherwise. In that case neither *true* nor *false* would be a true description of it. That is what we seem to be able to imagine about bravery. So unless we know there is such a worldly factor, we cannot correctly call the statement true, false, or either true or false. Where semantics provides no blanket guarantee—as here it seems not to—it is up to the world to furnish one; and this the world may fail to do.

Dummett's idea is that nothing in what a statement represents itself as doing, or anyway, in what some do, guarantees that such a situation could not arise. So, as he puts it, 'There will be no general guarantee that every statement is either true or false' (1978, p. xi). Other philosophers find themselves in Dummett's neighbourhood on this point too. I do, for example. Here is my idea. In the business of sublunary description, whatever, world willing, *might* count as representing things as they are, or as they are not, might, for specifiable ways the world might be, count as neither. So however a statement represented things, that would be, in specifiable circumstances, neither representing truly, nor representing falsely. In such circumstances it would neither be so that the statement was true, nor so that it was false. Whatever might be true, or false, might also demonstrably fail to be either.

All this, of course, is me, not Dummett. He and I agree that there is not in general (I say ever) a certain kind of guarantee. When it comes to what it means for there not to be, we part company. For I have denied his first intuition; I have read sense into what he declared senseless. So, like Williamson, there is anyway one problem I do not face. If there is any difficulty in reconciling Dummett's two intuitions, *that*, at least, is not a problem for me.

To put pieces together, Dummett has a position consisting of two points which can be summed up thus:

1. It is senseless (not true) to say of any statement that it is (or in such-and-such conditions would be) neither true nor false.

2. For some statements, it need not be so (there is no guarantee) that either they or their negations are true, that they are either true or false, or that they, disjoined with a negation of them, form a true statement.

Can one really say all that?

What is to block this line of thought? If it is senseless to say of *any* statement that it is neither true nor false, then, pick your favourite statement, it is senseless to say this of that one. It is, e.g., senseless to say this of a statement that Jones was brave. It is not senseless because the description 'neither true nor false' is *per se* senseless. That might, e.g., be a true description of my hat. What Dummett must mean is that this is senseless in the sense that it does not describe a possible way for a statement to be: no condition a statement might be in is correctly described as that. If Dummett is really in a position to say this, then he knows it. So *we* know it, or could come to know it. Dummett could prove it to us.<sup>2</sup> In particular, we know this about the statement that Jones was brave. We know—are guaranteed—that that statement is *not* in a condition of being neither true nor false. Again, we have, or could have, a proof that it is not—since Dummett can prove this of any statement. What better entitlement might we need to assert that that statement is either true or false? Must it not then *be* the one or the other? Could we be not entitled to assert what we know to be true? (There seemed no guarantee of something in the world to make S either true or false. We thought we could conceive that as contingent. But the appearance of contingency proved illusory.)<sup>3</sup>

There *might* be various ways of resisting the transitions here. Few seem promising. One might, though, try this move. In the train of thought, we need, at some point, to move from something like 'It is guaranteed not so (untrue) that S is neither true nor false' to 'It is guaranteed that S is either true or false'. Roughly, we seem to be moving from something of the form 'Not (not (S or not-S))' to something of the form 'S or not-S'. If we suppose that that 'or' and those 'not's are the stuff intuitionist logic speaks of, we may note that the inference here is not intuitionistically valid. So if we suppose that an argument is valid only if intuitionist logic, applied as just sketched, certifies it as valid, then the train of thought breaks down. On that assumption, Dummett may hold both of his intuitions, or so it seems.

Suppose there are overwhelmingly good reasons in favour of both of Dummett's intuitions. Suppose we may have them if, but only if, intuitionist logic,

<sup>2</sup> We must idealize away from stubbornness, etc.

<sup>3</sup> One might also argue in the reverse direction. For some statements there is, we are told, no guarantee that they are either true or false. (For that there would have to be something to make them so.) It is not necessarily true that they are the one, or are the other. So, one might think, possibly not. But if possibly not, then it cannot be senseless to describe a statement as neither true nor false. Perhaps we cannot know that it is that. But that certainly describes a condition that a statement might be in. If this line of thought breaks down, it will have to be in the same way as its reverse counterpart does.

applied as sketched, is the standard of correct inference. That might be a powerful argument for intuitionism. But, I will argue, the first supposition is incorrect. So there is, so far, no argument. There may be good reasons for preferring intuitionist to classical logic. But there are none in the area just scouted.

## 2. Understandings

Frege wrote:

The above-cited defects [of natural language] have their source in a certain plasticity and mutability of language, which, on the other hand, is a condition of its ability to develop and its many-sided usefulness. In this respect language can be compared with the hand, which, despite its ability to apply itself to various tasks, does not suffice. We create for ourselves artificial hands—tools for particular purposes—which work more precisely than a hand could. And how is this precision possible? Through just that rigidity, inflexibility of the parts, whose lack in the hand makes it suitable for so many different things. So, too, natural language does not suffice. We need a system of signs from which ambiguity is banned, the stricter logical form of whose content cannot slip away. (Frege 1882*b*, p. 52; trans. p. 86)

Our language has the advantage that it never, or rarely, leaves us at a loss for words. The descriptions it provides make the novel describable. To describe it is *ipso facto* to relate it to the known; generally in useful ways, ones that make the novel tractable. There is, as Frege notes, a price. Our language is not rigid. It does not foresee its own application in, or to, every circumstance in, or to, which it may prove to need to be applied. Occasions for applying it are left some role in determining how, on them, it is to be applied. That is a device that allows our thought to apply to the world, ensures sensible, and *true* things to think, no matter how the world might be. In another place Frege writes:

It is astonishing what language achieves, in its expressing, with a few syllables, unsurveyably many thoughts; that it immediately finds clothing for a thought that has just been grasped for the first time by an earthling, by which another can recognize what is entirely new. This would not be possible if we could not distinguish parts in a thought, which correspond to the parts of a sentence, so that the structure of the sentence can count as a picture of the structure of the thought. Admittedly we really use a simile when we carry over the relation of whole and part to thoughts. But the simile is so natural and entirely on target that we can regard its occasional lame bits as merely perturbations. (Frege 1923, p. 36; trans. p. 390)

The syntax of a natural language is productive. From a relatively small vocabulary it generates a vast stock of sentences—for any list of them, a novel one. There is a corresponding semantic structure to a language. Idioms aside, a whole expression has a meaning which is predictable from the meanings of its parts. So we have, with a given language, an (as good as) indefinitely large stock of vehicles for

expressing a thought. Each has a definite meaning, predictable from its parts, making it a description of a given way for things to be. But we miss something if we see this as the *whole* ground of our ability to express novel thoughts. There is also the plasticity Frege spoke of.

'Pigs grunt' is an English sentence. It provides a particular means for describing pigs. If, in speaking English, you want to describe pigs as grunTERS, these are just the words for you. Use them as meaning what they do, and that is what you will have done, or then counted as having done. Now, when would pigs be truly describable as grunTERS? Suppose that modern pigs are so well reared that none would ever dream of grunting (though they have lost none of their capacity). On one understanding of being a grunTER, pigs are still that. On another they are not. One *might* understand being a grunTER in either way. (Eastwood is not a grunTER *merely* on grounds that he could do it if he tried—at least on many occasions for classifying actors. GrunTERS are the ones you can depend on to do so regularly. But if we are casting, and the other actors in the pool—new men that they are—could *not* do it if they tried, then, for that purpose, Eastwood *might* count as a grunTER *merely* on those grounds.) Being a grunTER admits of understandings; there are different views of what would count as being that.

The English 'grunt(s)', as a description of things as grunTERS, is, in itself, no more tied to some one of these understandings than to any other. That means that there are various things to say—each on some occasion—in describing something as a grunTER (or doing what then so counted). For one may describe things as grunTERS where there is an understanding to be had of what, for purposes of *that* describing, would count as being that. There are truths to tell, but also falsehoods, in describing the imagined pigs, or Eastwood, as grunTERS.

Similarly, 'It's blue' is an English sentence which may be used, *inter alia*, to describe some ink as blue. That ink might be a substance which is black in fluid form, but writes blue. There is, again, an understanding of being blue on which such ink is blue, and an understanding on which it is not. There are, again, two corresponding things to say in calling ink blue in saying 'It's blue'. There is a further point that may be put in saying: understandings may admit of understandings. On one understanding of ink being blue, ink that writes blue is blue. What of ink that makes blue marks which then slowly disappear? On one understanding of writing blue it does so. So on one understanding it is blue on that understanding of being blue. On another it does not. So understood it is not blue on that understanding of ink being blue. There are various things one might say in describing ink as blue on an understanding on which writing blue counts as being blue. The understanding given English words, on a given speaking of them, bear may well choose between some understandings of being as those English words describe things, without choosing between all. So it may allow various understandings of being as those words as used said.

The moral. English furnishes descriptions that speak of given ways for things to be. The ways to be they speak of admit of understandings. So, in giving one of



those descriptions, one might say any of many things, all the while using given English words as meaning what they do. That is one way plasticity allows for expressing novel thoughts. We come to see—the world may teach us—novel understandings of being thus and so, where our language provides us one or more descriptions which are, specifically, descriptions of being *that*. For any understanding of a thing, or things, being that, we may, in so describing things, say them to be as they are if that on that understanding. We learn new things one might understand by being that; new things being that might be taken to be. We thereby gain the ability to use old descriptions (ones already at our disposal) to new ends; for expressing new thoughts.

There is a flip side to plasticity. It is a language's expressions—English's, say—that mean one thing or another. It is in, or by meaning what they do that they provide us with the descriptions for things that they do. What given words mean does not determine fully what it would be for things to be as one described them in using those words on some occasion. For their meaning makes them speak of what admits of understandings. A use of them, on an occasion, may bear an understanding that settles some of what their meaning on its own does not as to what would count as being as *thus* described. But there need be nothing in a particular use to settle all such questions. There are various understandings of being a grunter, so of being as the English 'grunts' says. There may also be various understandings of being as Pia said in using 'grunts', in saying, say, 'De Niro grunts'; competing understandings equally compatible with all there is to understand as to what it is she said. And where that is so, the world may conspire to make it matter. De Niro may count as being as she said on some of these understandings but not on others. The next task is to see how that matters to Dummett's first intuition.

### 3. Isostheneia

One might have the idea that, roughly, whatever might be true or false must be (or, more cagily, could not fail to be). More precisely, if it *would* be true, or false, of the way things were, if things were thus and so, then it *is* the one, or the other; or, again more cagily, it cannot fail to be the one or the other. My hat and Cleveland are neither true nor false. But then, they would not be no matter what, unless, somehow, we got them—in the right way—to represent things as so. Dummett and Williamson each have some version of this idea. Each takes it that there is a non-trivial way of demarcating the class of things that could not fail to have a truth value; and that it has to do with the sort of representing some things anyway do. For Williamson, there is a class of utterances that cannot fail to have a truth value. An utterance is in it if it is correctly described as having *said that* thus and so. For Dummett, statements form the relevant class, where a statement is what represents itself as aiming for, and achieving, a certain end: saying how things are.

The flipside of plasticity threatens this idea. A given description for a thing, as provided by some English phrase—a description of a thing as blue, say—admits of understandings. It describes things as a given way, where there are various (competing) understandings of being that way. A given thing may be that way on some of these, but not on others. If so, one may speak truth, but also may speak falsehood, in using that description to describe it. Suppose one uses that description to describe such a thing. That use of it will bear a certain understanding. It may be part of that understanding that one said the thing to be the way that description speaks of on such and such understanding of its being so. If the thing is some ink, say, black in the bottle, but writing blue, there is an understanding of its being blue on which it is that, and an understanding of that on which it is not. But, on the understanding one's use of the description bears, one may have said it to be blue on an understanding of that on which writing blue *is* being the way one said. Or, again, on one on which *merely* that is not being the way one said. For all that, there may be pairs of understandings of ink's being blue between which the understanding your words bear does not choose. Disappearing ink, which writes blue initially, but fades to invisibility, is, on one understanding of ink's being blue, blue. It contrasts with red disappearing ink, making it just the sort of thing a spy would want in certain situations. But on another understanding of ink's being blue, it is not blue; it is invisible. Your use of the description *may* bear, neither an understanding on which ink's being blue on the first of these understandings of its being so is its being as you said, nor one on which ink's merely being blue on that understanding, if it disappears, is its not being as you said. The understanding your words bear may simply not foresee, or speak to, such things as disappearing ink.

So far there is no harm. Indeed, the condition described is just that, it appears, of all our sublunary descriptions. But there is room for the world to make trouble should it so conspire. Suppose a statement spoke of things, or a thing, being a certain way; being thus and so. The world *may* then be arranged thus. There is an understanding one might have of things being that way on which that is how things are. There is, equally, a possible understanding of things being that way on which things are *not* that. It is not so that the statement is to be understood as saying things to be that way on some understanding of the one sort, rather than on some understanding of the other. No understanding the statement bears chooses between these: makes it speak of being the way in question on the one understanding, or on the other, of being that way. Things *are* the way they were said to be on one understanding. But if that is a reason for counting things as being as described, it cannot be a decisive one. For things are *not* the way they were said to be on another understanding one might quite well have of things being that way. Which is equally reason for counting things as being *not* as described. And vice versa. Asked whether someone is brave, we do sometimes say, 'Well, she was and she wasn't'. But whatever exactly that means, we cannot seriously count things as both ways at once. Our ways of describing the world

allow such situations to arise. What we have in such a case is a sort of natural *isostheneia*—natural because, in contrast to the use made of this notion by the Pyrrhonians, the point is not about our, or the, *justification* for thinking, or asserting, thus and so, or about how one might know such a thing. Natural *isostheneia* is not a source of ignorance. The point is that there is something which all the reason there is—all the facts that obtain—fails to decide. Everything about the way the thing described is makes it equally well, and no more, the sort of thing the description used describes (on the understanding on which it was used) as not that sort of thing at all.

What allows for the world to be arranged in such a way is the ways of representing it available to us, and the limits to the understandings our uses of these ways bear, and are capable of. We speak of ways to be that admit of understandings. The right understanding of our doing so (on an occasion) may fail to choose between some of these. The world may then arrange itself to correspond to some, but not others, of these understandings of being the way in question. There is then something we cannot do. We cannot use that description rightly of the world, but in a way that does not choose between these understandings. But whatever inability there is here is not, nor does it result from, ignorance.

To illustrate, I draw from Austin (1979/1950, p. 128). Zoë, in her wing chair before the fire, though in her prime, has entirely unexpectedly just expired. There is an understanding of being at home on which that is where she is—dead in her chair. There is another on which she is no longer at home since, as the euphemism suggests, the departed are no longer with us—are nowhere; Zoë, alas, is no more. Now suppose that, coincidentally, moments after Zoë's last breath, Max and Pia are discussing whether to pay a visit. Pia suspects Zoë may be out. But Max, ever confident, says, 'Oh don't worry. Zoë will be at home.' Are things as Max said? On the one understanding yes, on the other no. But recall that Zoë's demise will be a tremendous shock. No one anticipated it. So Max was not to be understood to speak of her being at home on an understanding on which the dead might be there, nor to do so on one on which the dead could not be. The understanding Max's words bore in no way addresses such issues. (That *is* the understanding they bear, not ignorance as to what it is.)

How should we describe a case of natural *isostheneia*? To avoid confusion, it is best to speak some metalanguage. It is not hard to see why that should be. Normally our aim in talking is simply to describe the world, or some specific bit of it. Our attention is on the world, not on our describing of it. We are not concerned to describe the relation between the descriptions we use and the world—both the descriptions and the world at once. That is our customary unreflective way of using language. The perspective we thus adopt takes for granted a stock of descriptions unreservedly at our disposal for just that purpose; that is, ones which (as we would use them) actually do, unproblematically, describe the way things are—or, again, are not. Just that is what is not so in a case of natural *isostheneia*. Here the relation between a description and the world

takes a turn unenvisioned from within this normal perspective. It has a third status which makes that perspective inoperable. The description turns out to be, as Austin would have said, inept. Its relation to the world can then only be seen rightly from, as it were, a self-conscious perspective, bringing the description itself into view along with the world it attempts to describe. Object language suggests (though not inevitably) what I have called the normal perspective. If we use it to describe natural *isostheneia*, our denials and negations will inevitably tend to be heard as spoken from within that perspective, with all that perspective presumes as to how the words we use fit the world. So hearing them will, in this case, create just the wrong impressions.

We must stand back and look at the description used and its particular way of fitting (or not) with the world. In a case of natural *isostheneia* the description was inept. As I will say, things are not correctly describable in its terms, either as being the way it speaks of, or as being not that way. It is not a way of saying what the facts are, nor, equally, are not. For there is no such fact as that things *are* the way it speaks of (such-and-such way, where that is the way it speaks of); nor, equally, as the fact that things are not that way. There is nothing in the way things are, nothing about the thing described, to make them, or it, the sort of thing that that description (as used) speaks, or spoke, of; to make it a fact that they are. Equally, there is nothing to make them *not* thus and so, where that is the way the description speaks, or spoke, of. That description, understood as it then was to be, is, *de facto*, not a right way of viewing either how they are or how they are not. For the way things are counts equally as fitting it and as not. Nothing in the way things are decides which it does. So it proves an inadequate means for describing things. It is thus not a way for us to talk; not a way we might use to specify what the facts are, or what is a fact.

As long as the way as statement speaks of things being is one that admits of understandings as to what being that way would be, and as long as the understanding *it* bears does not choose between every competing pair of these, that statement is susceptible to a condition of natural *isostheneia*, a third condition beside describing things as they are and describing them as they are not. Natural *isostheneia* is a way for the world to be, determinately, knowably, arranged. It is not an epistemic condition. The possibility of competing understandings of being whatever way we know to speak of things being explains why there should be that possibility, and how there could be.

#### 4. Tertium Datur

What should we say about truth in a case of natural *isostheneia*? Assume that Max's statement is such a case. Speaking truth is, one might think, 'telling it like it is', which, one might think, is saying things to be as they are. On one understanding of being the way Max said, things are that way. But, again, if that

is *some* reason for judging his words true, it cannot be decisive. For on another possible understanding of being that way, things are not that way, which counts equally against counting what he said as true. If speaking falsely is saying things to be as they are not, we get the same result as to whether his statement is false. So far *isostheneia* percolates upwards. If there is *isostheneia* for the question whether Zoe was at home (on the understanding of that on which Max said so), then there is equally *isostheneia* for the question whether Max's words were true. Neither (straight) answer to either question can count as the fact of the matter. And similarly for the question whether his statement is false.

Are there any further reasons in the nature of things that might tip the balance one way or the other? There is an understanding of being true (and likewise of being false) on which we expect truth to have a sort of transparency. To describe that phenomenon correctly, we need to resort again to metalanguage. The core idea is: if someone said that the way things are is thus and so, and did so truly, then I can say so too. In saying things to be just the way he did, I will state truth. The truth of what he said guarantees that. So from the truth of what he said I can draw a valid inference, the conclusion of which I can express in saying things to be precisely the way he did. I will describe things correctly in doing so. I will do no less than saying things to be the way they, in point of fact, are. It will be a fact that things are the way I say.

A parallel idea might hold for being false. If someone said things to be thus and so, and said so falsely, then I may state the truth in saying things not to be that way.<sup>4</sup> So, similarly, for the inference I may draw, and the way I may state its conclusion. If the way he said things to be is thus and so, then from the falsity of what he said I may conclude what I would state in saying things to be not that way. In saying that I would do no less than describing things as being a way they as a matter of fact are.

I think we clearly often do understand truth and falsity as transparent in this sense. If we do, then we must hold Max's statement not true, and, equally, not false. For to say Zoe to be at home on the understanding of that on which he did is not yet to describe things as in point of fact they are; nor, again, as in point of fact they are not. And we must say the same wherever *isostheneia* reigns. If *isostheneia* reigns for a given statement, it will reign for what I say if I say so too. If I say things to be just the way that statement did, it will not be a fact that things are as I just said; nor a fact that they are not. So such a statement does not satisfy the condition transparency imposes. On this understanding of truth and falsity, a statement for which *isostheneia* holds is not true, and not false. *Isostheneia* is thus a determinate *tertium* alongside those conditions. It consists in a distinction's having inescapably broken down in a particular case; whereas truth and falsity, on this understanding, consist in part in relevant distinctions actually classifying the way things are.

<sup>4</sup> Wittgenstein: 'If I say falsely that something is *red*, then, for all that, it isn't *red*' (1953, §429).

This understanding of truth and falsity does lose something. If we are told that it is not true that thus and so, or that someone said that thus and so, but that is not true, we often expect to be able to infer what we could express in denying precisely that things were *that way*. With *isostheneia*, and this understanding of truth, that inference does not always hold. In particular circumstances we may be able to see that *isostheneia* is outwith the range of what is countenanced by a given denial of truth. We may then be right to draw the inference. But there is, as it were, a cagier sort of denial, and a more subtle way for it to be correct. It might be true to deny the truth of something because for it *isostheneia* reigns. Unsurprisingly, *isostheneia* means, where it arises, that we cannot carry on our ordinary mode of object-language talk, from our usual perspective, with complete nonchalance. Something has to go. This is a plausible candidate.

But is there an understanding of being true and being false on which we can still maintain that there is no way for a statement to be something other than either true or false? We have already seen such a way. We merely delete transparency from our expectations, holding on only to the schematic principle:

If what *S* said is that *P*, then *S* is true iff *P*.

Then if *isostheneia* holds for *S*, it holds equally for the statements that *S* is true, that it is not true, that it is false, and that it is not false. We have no way of talking about *S* correctly; no description we can use to describe it rightly. But, for that price, we also have no description of a status alongside truth and falsity; or at least not one on which a statement might be not true, but, for all that, not false. Perhaps it all depends on whether one would want such a thing.

But is there no understanding of truth and falsity on which Max's statement comes out either true or false—preferably, in that case, false? It is clear what we would need to have one. We would need a disjunctive conception of falsity (with a consequent lack of symmetry between truth and falsity). Let truth be transparent as above. Now understand falsity so that there are two ways of being false: either things are not as they were said to be; or *isostheneia* reigns. (In the long run this will not save things, for a reason to be mentioned at the end of this section: *any* scheme for dichotomizing the world is liable to break down. But neglect that for the moment.) If we understand falsity in this way, then Max's statement is (determinately) false. I do not think we ever do understand being false thus. I see nothing to recommend this understanding. But I need not go to the wall over that for the present purpose. For if we understand Dummett's second intuition in terms of *isostheneia*, that will prove enough of a *tertium*, even if only at object level, to block any move from his intuitions to intuitionism.

Dummett's idea was that if (*per impossibile*) a statement could be not true, but not false, that would have to be through some special understanding it bore according to which, under such-and-such circumstances, it was so to count. A statement bears an understanding as to when—on what condition, under what conditions—it would be true. What it would need, the thought is, to

be neither true nor false, is to bear an understanding as to when—under what conditions—it would do this. Dummett quite rightly observes that statements bear no such understanding.

But that is not how truth-valuelessness arises. Take any pair of categories designed to be mutually exclusive and exhaustive (over some specified domain): if things are a certain way, they go in the one category; if they are not that way, they go in the other. No matter what that way, it is always conceivable that the world should cast up items that are unclassifiable within that scheme: the way those items are does not make them the sort of thing to fall into either category. For example, understand what you will by being blue. There are conceivable items that count neither as blue, nor as not, on that understanding of being so coloured. And, I argued, that possibility inevitably infects our talk of truth and falsity. If a statement said things to be thus and so, then if things are that way it is true, and if they are not it is false. But things may fail to classify either as that way or as not.

But how, one might ask, can that be? If a thing is not correctly counted as in the first (positive) category of the pair, must that not mean that it is *not* the way in question, so counts as in the second category? Natural *isostheneia* shows how that need not be so. An item may fail to be correctly classifiable as in the positive member of the pair, not because it is *not* the way in question—as if that were a way of saying what is so—but rather because there is, in nature, no more, but then again, no less, reason for placing it there than for refusing to. The way it is makes it as much the sort of thing that belongs in the one category as the sort that belongs in the other.

If an item has this status with respect to a given pair of categories, we can replace these with another pair into which it unproblematically fits. So, for example, if Max's statement counts neither as true nor as false, on our usual understanding of being true and being false, we can find another understanding on which statements like Max's, for which *isostheneia* reigns, are to be classified as—say—false. On that conception, Max's statement is—of course—unproblematically false. What we cannot do, though, is to define a pair of categories in a way that is guaranteed to eliminate *all* possibility of items which, on grounds of *isostheneia*, are unclassifiable in that way. Statements neither true nor false require no *special* understanding of what they are about to make them that.

## 5. Logic

If the question is what logic is right, the crucial point in the present discussion may *seem* this. Take an arbitrary collection of statements, for each of which there is a way which is the way it said things to be. Is it guaranteed that every member of the collection is either true or false? Does that follow merely from those statements being the ones they are? Or from each having the feature that

there is such a thing as how it said things to be? My answer is no. For some collections, we may know that their members are all either true or false. But, at least if the statements describe sublunary affairs, that does not follow merely from their being the statements they are. On one understanding of truth and falsity, some such collections contain items that are neither true nor false. On another, some contain items for which *isostheneia* holds on the questions whether they are true and whether they are false. Nothing makes it a fact that they are. Nothing makes it a fact that they are not.

The question may seem crucial because it *may* seem a touchstone for classical logic. Three questions arise. Have I denied bivalence? Have I said what, if so, means that classical logic is incorrect? Have I said what militates in favour of sublunary intuitionism? My answers are as follows. On the first question, I do not know. It all depends on what you mean by bivalence. On the second and third, no.

What bivalence might be is a story too long to tell here. Two brief remarks. First, it is part of what stating is that to state is to purport to contribute to collections of statements for which bivalence holds (hence to which classical logic applies). To state is to purport to say how things in fact are; *a fortiori* how they are or, again, are not. Second, as long as we see statements from an object-level, internal perspective—as long as we view them as saying things we might say, or deny, ourselves, what we agree or disagree with (or do not yet know about)—we see them in a way that assumes, or presupposes, bivalence. We are treating them as what is either true or false; so as subject to the principles of classical logic. Our ordinary way of using language to describe the world, focusing on the world, and not on ourselves describing it, presumes bivalence. Taking it to hold is the correct attitude in assuming this perspective. Or, better, it is just part of what it is to assume it. Of course, there is no law that says that presuppositions must be true. Our ordinary stance sometimes breaks down. The descriptions we use sometimes prove inept.

The second question. There are questions for which *isostheneia* holds. Was Zoë at home, on the understanding of being there on which Max said so? Correspondingly (let us so understand truth) there are statements that are neither true nor false. Does that mean that classical logic is wrong? That surely depends on what classical logic says. So let us see what it does say, in the simplest case—sentential logic.

So restricted, classical logic is the theory of certain forms there are for a statement to take: truth-functional forms. The truth value of a whole of such a form is a function of the truth values of specified parts of it (a truth function being a function from  $n$ -tuples of the values true and false to those values). Logic tells us ways various such forms relate to others. Notably, if instances of certain forms are true, it tells us, then a related instance of a further form, with specified common constituents, is.

In speaking of such forms, classical logic says something about any collection of statements which are all truth-valued—thus which contribute to arguments



for functions in terms of which the forms it speaks might be defined. Suppose some members of such a collection are of forms of which it speaks. So the constituents of those instances of those forms will also be truth-valued. Then classical logic may say various things about how those truth values relate; for example, if such-and-such of those statements are true, so are such-and-such others.

There are things no logic could say. It is not a law of any logic that such-and-such specimens of spoken English are, in fact, of forms of which classical logic speaks. Nor is it a logical law that English contains *any* truth-functional connectives. Nor does classical logic tell us that such-and-such specimens of English are truth-valued. How could a logic tell us that?

Suppose there is a language—let's call it *English*. Suppose it contains a sentential operator spelled 'not'. Ignoring the syntax of English, let us write the result of applying this operator to a sentence, 'S', 'not-S'. Suppose that 'not' is not truth-functional. Then 'not-S' is not of any form of which classical logic speaks. Whereof one does not speak, one cannot be mistaken. Suppose, for example, that, whatever form 'not' has, from what was stated in some 'not-not-S' it is not always correct to conclude what would be stated in the corresponding 'S'. Well, as the French are wont to say, 'Tant pis!' Classical logic is none the worse for that.

Classical logic is interesting, insofar as it is, because, often enough, we are entitled to auxiliary assumptions, given which we can, for certain purposes, correctly view certain statements as of the forms of which it treats. Given the assumptions, we can then consult the logic to see what inferences to draw—what inferential relations hold between those statements when they are so viewed. About *that* classical logic is not wrong. (Whether the assumptions are justified may be another matter.)

Let us approach the third question. If statements may be neither true nor false (not to put too fine a point on it), that opens up new possibilities for understandings a connective may bear. Conspicuously, it opens up new possibilities for understanding negation, some of which we certainly do exploit. Could there not be other, non-classical, logics which treated in forms that were sensitive to some of the features connectives would have on some of these understandings? Might some of these other logics not thereby be more useful than classical logic in informing us which inferences it is actually right to draw—the inferential relations that hold between the statements we actually make? Might not intuitionist logic be pre-eminent among these?

The first two of these questions merit a definite 'Maybe'. The third is in a way misguided, since the issue was meant to be not whether intuitionist logic might be useful for some purpose or other, but whether, in some sense of 'correct', it, and not classical logic, is correct. That is certainly what Dummett has a thesis about. In any event, the answer to the third question is that it does not seem so.

Let us consider negation, as expressed, for example, in using 'not'. On one understanding of negation, all that it does is to flip truth values: a statement

'not-S' is true iff the corresponding 'S' is false, and false iff that 'S' is true. On object level, 'not-S' is true just in case the fact is that things are not that way which is the way 'S' spoke of their being. That is the understanding of negation on which *isostheneia* is a matter of things being, neither such-and-such way, nor not that way; on which it is neither so that they are that way, nor that they are not. And it is one on which *isostheneia* holds for 'S' iff it holds for 'not-S'. But there are other understandings of negation. We can, and sometimes do, understand negation in this way: 'not-S' is false if S is true, but true if either S is false or *isostheneia* reigns. That is the most usually discussed form of so-called external negation, evident in uses of 'not' with modifiers: 'Zoë was not quite (exactly) at home; but not quite *not* at home either.' (Note the mixed understandings of 'not' in such a case. Such is the way we in fact think.) We also sometimes understand negation this way: 'not-S' is false if S is either true or false, but true if S is neither. We might think of this as *strongly* external negation. I prefer to call it *Lombardi negation*, since it is evident in that great coach's reputed remark, 'Winning isn't everything; it's the only thing.' All these are possible understandings of negation.

The point about all such understandings of negation, though, is that double negation elimination holds for all of them, as is easily checked. So (as one would expect from the way I explained them) none pushes us at all in the direction of intuitionist logic.

What, then, *would* move us in that direction? Let us return to Dummett's example of uninvincibly brave (or lack of it). Dummett's favourite locution for describing the case is this: we are not *entitled to assert* that Jones was brave, or that Jones was not brave (see 1959, p. 16). (Nor that Jones was either brave or not brave.) We may irredeemably lack such entitlement. However we understand 'entitled to assert', it must express a notion on which Dummett needs to place a great deal of weight. I have explained various understandings of negation in this way. I have divided ways the world may be arranged into cases. Things may be thus and so; they may be not thus and so; *isostheneia* may reign. I have then explained an understanding of negation by explaining its results for each of these cases.<sup>5</sup> As long as explanations run that way, the result will not be intuitionism. If intuitionist logic is to be the *correct* logic for our thought, such explanation must be illegitimate. We must not be able to explain understandings, of connectives, at least, in terms of how the world is anyway, but rather *only* in terms of what we are entitled to say. Dummett may then try to work out a well-motivated account of the distribution and transmission of entitlement that makes intuitionist logic the correct way of representing the relations between what we are entitled to say.

<sup>5</sup> I did not, and need not, suppose that every conceivable arrangement of the world would, or must, classify, determinately, as falling in one or another of the categories I thus set up. That would be to repeat what I regard as a mistaken view as to whether there may be schemes for dichotomizing which are guaranteed, just by being the schemes they are, actually to dichotomize.

Why think an account of entitlement relations would validate intuitionism? One reason one might have is that it must make good on Dummett's second intuition. For which, Dummett takes it, it must leave the following situation possible. Given an arbitrary amount of information we might have, we are not entitled to assert that Jones was brave. Likewise, we are not entitled to assert that Jones was not brave. And, again, we are not entitled to assert that Jones either was brave or was not. The thought is: an account of entitlement that respected that restriction would make intuitionism right.

Dummett's second intuition, as he repeatedly insists, rests on this idea: a statement is true, or, *casu quo* false, only if there is something in the world in virtue of which it is true (*casu quo* false). Let us ask a question. Why should we not be able to be in this position? We know enough about bravery, and about Jones, to *know* that nothing about him makes him the sort of person correctly described as brave, and nothing about him makes him the sort correctly described as not brave. We *know* that neither description is correct. We should, then, be entitled to say so. And a way of saying so would be: it is neither true that he was brave, nor true that he was not. In material mode (as we naturally do put such things) he was neither brave nor not brave.

An answer had better not be: one could never know such a thing—one's evidence would always be too flimsy. As if something might always still turn up to show that Jones was brave, or to show that he was not. It is conceivable that something should turn up to show that Paris is not in France. That is conceivable; but, for all that, we know that it is not so. We know that Paris is in France; nothing will turn up to show otherwise. Similarly—if flimsy evidence is the only problem—we may *know* that there is nothing to make Jones count as brave, and also nothing to make him count as not brave. Nothing *will* turn up.

An answer must not rest on bad epistemology. It must appeal to some internal feature of entitlement. The idea would have to be that I have described a logically impossible situation. That idea would be spelled out this way: if we had proof that nothing made Jones count as brave—a proof that there can be no proof that he is, that that is unprovable—then that would *ipso facto* be a proof that Jones was not brave. That is a familiar intuitionist idea, plausible enough for at least some of mathematics. But as an idea about sublunary thought it is just a misunderstanding. That is not because something might be unprovably so. That is as may be. It is rather that a certain condition might, even *provably*, obtain. There may be provably no proof of P because *isostheneia* holds as to whether P. In the present case that possibility comes to this. Nothing in any proper understanding of being brave—of what bravery is—chooses between counting Jones as among the brave, and counting him as among those not brave. If he absolutely *must be* classified, one way of doing it is as good as, and no better than, the other. Which shows that there *is* no right way of classifying him—a fact about him we can perfectly well know. Which shows that it is not so that he must be classified at all. As soon as we recognize the possibility of *isostheneia*, and

understand cases such as Jones's accordingly, there is, as I have shown, no longer any push, at least in Dummett's intuitions, towards sublunary intuitionism.

## 6. Shadows

Dummett's break with traditional philosophical ways of thinking is radical. One question is why he breaks with it where he does rather than reject his first intuition, say, along the lines suggested here. Another is precisely where he does break with it. A full answer to the first question is at the very least beyond the scope of this essay. But we are in a position to locate his break with the tradition in a way that may be helpful. We can begin with a view of the break he will not make.

Some things have truth values. Some do not. What distinguishes the one sort of thing from the other? To have a truth value, an item must represent things in some particular way. That is the trouble with my hat. That idea is compatible with this one. A way of representing things may be such that, world willing, what so represented would have a truth value, but, world unwilling, what so represented would not. That amounts to rejecting the first intuition. One might think to stave off that idea with this one: there are ways for words, or other items, to represent things such that whatever *so* represented things could not lack a truth value, no matter how the world was. One might then add: whatever is, or could be, truth-valued represents in some such way. To hold Dummett's first intuition one must think that, as both Dummett and Williamson clearly do.

My main brief has been to argue, against both these last claims, but, most crucially, for present purposes, against the second. I have tried to show that all the ways some words count as having represented things admit of understandings. (Against the first claim I hold that all the ways there are *for* words to represent things do that.) For all the facts there are as to how given words represented things, the world may confront us with a situation which, on one understanding of being that way, is as those words represented things, but, on another possible one, is not. Then, by the argument of section 2, such words can count neither as true nor as false.

Ways for words to represent admit of understandings. So they require cooperation from the world if the understandings they might have are not to matter to whether they are true so as to leave them unevaluable as to truth. The core idea here—that understandings admit of understandings—amounts to rejection of a larger picture. It is the picture Wittgenstein rejects in rejecting a conception of propositions as, as he puts it, shadows.<sup>6</sup> The idea of a shadow is that of a specifiable form, or way, for words, or other things, to represent, which

<sup>6</sup> See Moore 1959, esp. pp. 260–1.

determines all that is determined as to when what so represented would be true (or false). If words represented in some such way, then everything that was so as to when, or of what, they would be true, or false, would follow from their having so represented (and the world being as it is). So the defining feature of a shadow is this: a shadow is a specifiable way of representing the way things are which, first, could be, or have been, representing truly, or representing falsely, and second is such that if anything which represents things in that way has a given truth value then, necessarily, everything which represents in that way shares that value. According to Wittgenstein, and to me, there is no such way for words, or anything, to represent.

Though it is not so by definition, it is plausible, if we reason classically, that, of necessity, a shadow's way of representing would be, no matter how things were, either a way of representing truly, or a way of representing falsely. To represent as a shadow does, the idea is, must be, either to represent things as they are, or to represent things as they are not; so what represents as a shadow does *must* have a truth value. For suppose, of some shadow, not. Then, if the world were right, the fact of representing as that shadow does would decide neither that one, or an item, represented truly, nor that one, or it, represented falsely. But then it is consistent with representing in that way, where the world was that way, that one might have represented truly, given further facts about how one thus represented. The shadow's way of representing does not rule out that possibility. Similarly for so representing in representing falsely. But then that supposed shadow is no shadow. In effect, its way of representing admits of understandings.

By this line of thought, if our statements are shadowed, then they must be either true or false. They are guaranteed to have either the one or the other truth value. But that result is one that Dummett cannot accept. Rejection comes hard for him, though, since he retains the idea that our statements, and our thoughts, are shadowed. That is why he cannot accept the way I have offered for a statement, knowably, to count neither as true nor as false. His adherence to the idea of shadows also shows in another pillar of his picture: the idea that anything that is so as to when our words would be true, when false, must be entailed by facts about some specifiable meaning—that is, representational form—that we have succeeded in conferring on them. So Dummett's rejection must consist in rejecting the line of thought leading from the idea of shadows to the idea that every shadowed statement must be either true or false. It is here that sublunary intuitionism serves well.

Why must Dummett think of our statements as shadowed? Here I must indulge in some degree of speculation. Perhaps it is because unshadowed thought threatens to be mind-dependent, or at least human-dependent thought. If our statements, and thoughts, are of no form that is a shadow, then, it seems, they represent as they do only given our reactions to their ways of representing, and

their roles in our lives. Unconjuring that seeming menace is a larger project than this essay, which I leave entirely for other occasions.

Whatever its sources, Dummett's break with the tradition may well be right. Perhaps we must stop thinking of the world in ways that make the semantic basis of classical logic make sense. That idea is too deep and pregnant to dismiss merely on account of anything said here. With the idea that our thought is shadowless, though, it becomes harder to make Dummett's case. For the idea is no longer required by much of what otherwise might.

# 6

## In insensitive Semantics

What is insensitive semantics (also semantic minimalism, henceforth SM)? That will need to emerge, if at all, from the authors' (henceforth C&L<sup>1</sup>) objections to what they see as their opponents. They signal two main opponents: moderate contextualists (henceforth MCs) and radical contextualists (henceforth RCs). I am signalled as a main RC. I will thus henceforth represent that position *in propria persona*. In most general lines the story is this: MC collapses into RC; RC is incoherent, or inconsistent, on various counts; SM is thus the only game in town. As to all of that, we shall see.

C&L's MCs are, in essence, followers of Carnap. Carnap's idea was this: if an expression's contributions to what is said in speaking it vary from speaking to speaking, find a variable with different values for those different speakings; postulate a (tacit) device in the expression, referring, on a speaking, to the value of that variable for that speaking, find a function from that value to the contribution of the expression, on that speaking, to what was then said. In that way a sentence might be thought of as a function from indices (n-tuples of values of all such variables contained in it) to particular things there are to say. The sentence varies its truth conditions from one speaking of it to another; those things there are to say do not. MCs see more such expressions with variable contributions than, let us say, was once thought. So they postulate larger indices than was once thought necessary.

C&L's RCs do, in some sense, more of the same of what MCs do. They certainly see more variable contributions of sentential constituents. In fact they see such variation across speakings (henceforth occasion-sensitivity) as, in principle, ubiquitous. They are not thus led to posit (larger) indices. If occasion-sensitivity is ubiquitous, it is obviously not to be domesticated by such a strategy: one cannot pass from items with different conditions on truth for different occurrences of them to items not eligible thus to vary by, so to speak, letting one's indices grow. For an RC, Carnapian elimination of occasion-sensitivity is just not on. Do C&L recognize that? Sometimes, perhaps. Other times, not.

<sup>1</sup> Cappelen and Lepore (2005). Page references in this chapter are to this volume.

## 1. Credo

Here is a succinct SM credo:

The idea motivating Semantic Minimalism is simple and obvious. 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. (p. 143)

There are languages. Languages have words. Words combine into complex expressions and sentences. The semantic values of words contribute to the semantic values of the complex expressions and sentences of which they are a part. Semantics is about how best to specify the semantic value of the lexical items and their contributions to the semantic values of complex expressions and sentences in which they occur. (p. 58)

Semanticists . . . agree that semantics is a discipline that aims to characterize systematically certain features of linguistic expressions and to do so in a way that captures general truths about languages, and not just truths about particular speakers in specific contexts. Characterizations of what speakers say . . . on the other hand, . . . aim to convey something about a particular act in a particular context . . . (p. 58)

What part of this do RCs disagree with? Answer: no part. Semantics is concerned with properties that expressions of a language such as English *have*. It is not *per se* concerned with properties that only some occurrences of them—if any—have. SMs and RCs disagree, as will emerge, over what properties expressions of a language *do* have. So let us keep the above ground rule firmly in mind as we proceed.

Truth is the central issue here. RCs deny that (e.g., English) sentences are in the business of being true or false. They are not (RCs hold) because there is, systematically, no such thing as ‘that which a sentence says to be so’. So (RCs) claim, there will not be truths to be told of this form: ‘The English sentence “Pigs grunt” is true iff pigs grunt.’ So that the idea of indices does not raise its head here, I will put this in terms of open sentences. What the English ‘\_\_\_\_\_ grunts’, or any other open English sentence, *means* leaves it open to say any of indefinitely many different things, at a time, of a given item, in using that open sentence of it. So there are no truths to be told of the form ‘The English “\_\_\_\_\_ grunts” is true of an item just in case that item is blah (blahs)’. Open English sentences do not, in that sense, have satisfaction conditions. These are, so far, claims, not a case. Reasons for them will be forthcoming.

C&L assign truth precisely that role in semantics which RC denies it. They insist that (unproblematic) declarative sentences—‘Pigs grunt’, ‘There are French women in Chicago’, ‘The oboe is double-reeded’—and even more problematic ones—‘Giraffes are tall’, ‘Sid has had enough’—express, *as such*, ‘minimal propositions’. Whatever else a minimal proposition is, it is (for them) something



truth-evaluable, and, as a rule, truth-valued. That, then, is the main issue between SM and RC.

Do C&L recognize that fact? Sometimes. They say, e.g., ‘That there is a proposition semantically expressed is presupposed by any coherent account of linguistic communication, i.e., accounts which fail to recognize a semantically expressed proposition . . . are incoherent’ (p. 144), whereas, they also tell us, it is a central tenet of RC that ‘[n]o English sentence S ever semantically expresses a proposition’ (p. 6). Here they see what the issue is. Other times, though, they seem to think things turn on acceptance or rejection of something like the credo. For example, they accuse RCs (and me specifically) of holding on to something which they call ‘the mistaken assumption’, of which they give two versions:

MA: A theory of semantic content is adequate just in case it accounts for all or most of the intuitions speakers have about speech act contents, i.e., intuitions about what speakers say, assert, claim and state by uttering sentences. (p. 53)

MA\*: If CSAs of the kind presented so far trigger the intuition that proposition *p* is said, claimed, stated, or asserted by an utterance *u* of sentence *S* in language *L*, then an adequate semantic theory for *L* should assign *p* as the semantic content of *u* (p. 54). (‘CSAs’ are ‘context-shifting arguments’, about which more later.)

But RCs think nothing of the sort. For an RC, semantics is emphatically not in the business of predicting what proposition would be expressed in some given utterance of a sentence. Nor do RCs think such things *are* predictable (as a function of some set of parameters). Exactly not. That would be a Carnapian approach—the sort of thing elaborated by Richard Montague, or David Kaplan. To an RC, such an idea simply fails to see what the phenomenon of occasion-sensitivity in fact is.

Incredibly, C&L accuse me specifically of holding MA. I cannot imagine why. Perhaps it is because I think that if a *sentence* were true under such-and-such conditions, and you spoke it where those conditions obtained, you would say something true. That sounds about right. But of course I do not think that, in that sense, any sentence does have a truth condition.

## 2. Argument

The core thesis of RC is that any way for things to be which an English (or etc.) open sentence speaks of admits of understandings as to when something would be that way. Any of many different things may thus be said of a given item in saying it to be that way. The same variety of different things may thus be said of it in using that open sentence of it. (Illustrations are about to come.) How might one argue for this?

C&L are fascinated by one style of argument. They call it a context-shifting argument (henceforth CSA). The argument begins with a contrasting pair. Each

member of a pair is a speaking of a sentence on which all its component expressions mean what they do mean, and such that the same items are spoken of each time (as being whatever way is spoken of by the open sentence they complete). The contrast lies in this: one member of the pair says something true (of those items), the other something false. The conclusion will be that that sentence may, while meaning what it does, say things to be any of various ways. It would do so on different speakings of it. It would thus have different truth conditions on different such speakings, some of these perhaps met while others were not, no one of which is the condition for the truth of the sentence as such. So what a sentence means determines neither when *it* would be true (there being no condition for that), nor when any speaking of it would. (Counter, *nota bene*, to what would be so by MA.)

Take, for example, the English 'Sid grunts' (spoken of Sid at a time). Sid has been in an accident. The fear is that he has lost his ability to grunt. Pia gives him a sharp blow in the solar plexus. He responds with a grunt. Which Pia reports (correctly) by saying, 'Sid grunts!' By contrast, Pia and Zoë are planning a soirée. Pia would prefer Sid not to come. (He knows too much.) So she says, 'Don't invite Sid. Sid grunts.' But Sid is the most urbane of men. He has no such habit. Pia speaks mendaciously, so falsely.

To establish RC in this way, one would need to make very sure that each contrasting pair was all it was meant to be, and to produce enough such pairs to produce conviction that the phenomenon is pervasive throughout English (etc.). Happily I need not do that here. C&L are convinced of the pervasiveness of contrasting pairs. It is not that over which they and RCs disagree. I merely reiterate an agreed ground rule. Semantic theory should not assign an expression a property *it* does not have; hence not one had by one member, but not another, of a contrasting pair. Such pairs may thus reveal some things semantics does *not* deal in.

C&L are also fascinated (for reasons to emerge) by what *they* would say sitting in their favourite coffee shop, speaking a sentence occurring in some contrasting pair. On this I think they mistake the force of RC, and of CSAs. Where there are contrasting pairs, the circumstances of a speaking must contribute *substantially* to fixing what is thus said, if anything even possibly true or false is to be said at all. Not just any circumstances will do the required work. So, taking Sid for granted, what will C&L say, in their coffee shop, in saying, 'Sid grunts'? Quite likely *nothing* either true or false. Suppose Lepore says, 'I know Pia said truly (the first time) that Sid grunts, and falsely (the second time) that he does. But does Sid grunt?' Then, probably, he has asked a question with no straight answer. Not enough is fixed as to how to take that occurrence of 'grunt' to make either 'Yes' or 'No' correct. Perhaps what C&L mean to ask is: 'Perhaps Sid grunts on one understanding of doing that, and not on another. But does he really *grunt*, that is, independent of any contingent understanding of what it would be to do that?' But if RC is correct, that is a bad question. Given RC, there is not, in addition to

all the facts as to whether Sid grunts on this or that understanding of his doing so some further fact as to whether he 'really grunts'. Grunting on one or another understanding of doing so is the most we ever do.

There is another way of arguing for RC. We begin with the plausible assumption that (the English) 'grunts' speaks of being a grunter, 'is blue' speaks of being (coloured) blue, and so on. We then ask ourselves such questions as, 'When *would* something be coloured blue? How about Lac Lemman? It has a blue appearance on this sunny day. But its water is not blue in the way that Lake Louise's water is green. Is it blue?' We find that one *can* understand (a lake's) being blue such that Lac Lemman is that way, but also such that Lac Lemman is not. Neither understanding is, so far as we can see, either required or excluded by what being blue is *per se*. So neither is required or excluded by 'is blue' meaning what it does. (So Lac Lemman can neither satisfy, or fail to satisfy, that English predicate as such.) We then observe that on some occasions for speaking of the colour of a lake, one *would* understand its being blue in the one of these ways, on others in the other. Since being blue is what 'is blue' speaks of, in using those words of a lake on the first sort of occasion one would speak of its being blue; hence of its being blue on that first sort of understanding of its so being. What one would thus say to be so, that is, would be so just in case the lake is blue on that understanding. Since what 'is blue' speaks of in meaning what it does is simply being blue, it is thereby eligible to speak of that, on occasion, on any of the understandings that being blue admits of. Indeed, if it were not—if it were somehow reserved, in point of meaning, for speaking of being blue on only certain understandings of being blue—then it would not speak of *being blue*. It would rather speak of being blue on a special sort of understanding of that. The meaning of an English expression makes it *for* saying a certain sort of thing in speaking English. The meaning of 'is blue', in making those words speak of being blue, makes it *for* (e.g.) calling something blue on the understanding there would then be of its being so.

Again the ground rule. We must not assign the English 'grunts', or 'is blue', a property it does not have. So we must not assign it a property in having which it would speak of something other than being blue. So we must not assign it a property which would make it *for* something other than speaking of what, on the occasion of the speaking, being blue would be understood to be. This rules out assigning 'being blue', or 'grunts', a satisfaction condition, at least if such a condition determines an extension. I will return to that point.

### 3. Reduction

I said it was a plausible supposition that 'is blue' speaks of being blue and 'grunts' speaks of being a grunter; that, in this way, open English sentences speak of particular ways for a thing to be. Suppose one asked, 'Just what *is* this thing,

being coloured blue? What way for a thing to be is that?' A neat answer might be: 'It is that way which a thing is iff C', where C stated a condition which was met by (and only by) some determinate class of things. Those things would then be 'the blue things' *sans phrase*. A similar answer for being a grunter would entitle us similarly to speak of 'those things which grunter'. And so on. Such an answer would be, in effect, a material mode version of a satisfaction condition for a predicate. But I have denied, on RC's behalf, that there is any such neat answer. So I have denied that there is any such thing as 'the class of blue things', or 'the class of grunters'—except, perhaps, where 'blue things', or 'grunters', bears some particular understanding. No class of things is the class of blue things *sans phrase*.

What I have said instead is that being blue is a way which a given thing may count as being, when one understands so being in a certain way (and on occasions where it would be so understood), and similarly count, on some understandings and occasions, as not being. In that sense, what sometimes counts as being it may sometimes not. Which may raise the question, 'When is it *that* way that is spoken of (or otherwise in question)?' To which my only answer is that if you are one of us (say, a speaker of some language enough like English), then that is something you are prepared to recognize. If you are not, then I have no formula—as it were, a cognitive prosthetic—to offer you with which to replace such competence. (Though there is much more to say on this topic, I do not believe there is such a prosthetic.) Which may leave one yearning for the neat answer I am unprepared to give. Which would show, in turn, what such a neat answer amounts to: a sort of reduction of being blue, or being a grunter, to something else—of speaking of such things to a condition on speaking truly of a thing in given words which actually decides when one would do so. I see no more hope for such a reduction than there is generally for reductions of anything to anything else. Read that way, Bishop Butler had the right idea.

Perhaps it is because C&L are so wedded to neat answers, or at least to such things as 'the class of grunters' *sans phrase*, that they cannot help misquoting me. They report me as saying, 'Since what sometimes counts as green may sometimes not be . . .' (p. 132). That 'be' is not mine. What I said, and meant to say, is: what sometimes counts as green may sometimes not *so count*. But they seem to think that I *must* have spoken of something counting as green *even though it is not*. Or else I just slipped. For they say:

The first claim made in the Travis quote above [the misquote] simply says that 'What sometimes counts as green may sometimes not be', i.e., a thing can satisfy 'is green' even though it is not green. No Radical Contextualist can deny this, because to do so is, in effect, to endorse Semantic Minimalism. It is constitutive of Semantic Minimalism that only something green can satisfy 'green', and so anyone who endorses this claim endorses Semantic Minimalism. (p. 136)

But what is meant to be constitutive of semantic minimalism here is what could only be said coherently at all if what I have called a neat answer were available

for the phenomenon of being green. By RC, you can only speak truly in calling something green if it *is* green on that understanding of a thing's so being on which you speak. So you can only speak truly in calling something green on occasions on which you would thus be calling it green on an understanding on which it is. There is no way in which something which *is* green on the understanding on which it was said to be may, for all that, really not be green, unless, irrelevantly, that is for it not to be green on some other understanding on which it was not said to be. There is no other sort of sense to be given to the expression 'only something green' in 'only something green can satisfy "green"' on which there is a coherent requirement here which, by RC, is flouted. Nor can semantic minimalism be defined by requirements not there to be imposed. (Though it might be defined by its trying to.)

Rejecting neat answers here does leave us with two different understandings of the notion *saying that*. To head off future confusion, I will sketch them. In using 'grunts' as meaning what it does, one speaks of (something's) being a grunter. If, on an occasion, Pia spoke of being a grunter, and said Sid to be that way, she did what it would then be to say that Sid was a grunter. On one understanding of *saying that*, for her to have done that is for her to have said that Sid is a grunter. If I say her to have said that Sid is a grunter, I *may* be speaking on that understanding. In which case I will have told you nothing about the understanding of being a grunter on which she said that. In which case I will have told you nothing which settles when what she thus said would be true. By compensation, I will have told you what can be known to be so merely in knowing that she used given words ('Sid grunts', say) as meaning what they do. On the other understanding, when I say Pia to have said that Sid is a grunter, I thereby identify what it is she said to be so. To do that I will have to have spoken on some particular understanding of being a grunter. So this understanding of *saying that* presupposes that there is some such understanding on which I speak. All comes to grief if not. In so speaking, I may inform you as to when what Pia said would be true. I also say what does not follow merely from her having used given words (such as 'Sid grunts') as meaning what they do. Conflating these two conceptions can lead to the strangest superstitions, such as that there is truth to be told in saying, 'If Pia said "Sid grunts", using those words as meaning what they do, then what she said is true iff Sid grunts'. I hope we can now avoid such blunders in what follows.

#### 4. Complaints

C&L have three main complaints against RC: first, it postulates context sensitivity where what ought to be tests for it are failed; second, it makes communication impossible; third, it is internally inconsistent. I take these in order.

#### 4.1 Failing Tests

RC tells us that, since ‘grunts’ speaks of being a grunter, there are many things to be said of someone in saying ‘He grunts’. C&L want ‘grunts’ to pass three tests. First, it should block disquotation: there should be no automatic inference from ‘Pia said “The lake is blue”’ to ‘Pia said that the lake is blue’. Second, there should be a bar to ‘collective description’. In C&L’s words, ‘From there being contexts of utterance in which “A *v-s*” and [ones in which] “B *v-s*” are true, it doesn’t *follow* that there is a true utterance of “A and B both *v*”’ (p. 99). Third, it should be possible to say (truly), ‘There can be false utterances of “S” even though S.’ While speaking truth in saying ‘S’, we should still be able to recognize, explicitly, circumstances in which, in saying ‘S’, one would speak falsehood.

These tests, applied with an attentive ear, *are* passed by ‘grunts’. The first first. Max, testing Sid’s medical condition, hits him sharply in the solar plexus. Sid grunts. Max reports the result to Pia in saying ‘Sid grunts’. At another place and time Zoë says to Pia, ‘Let’s invite Sid’. Pia replies, ‘Let’s not. Max said that he grunts.’ A false and deceitful response, to which Zoë replies truly (and indignantly), ‘He does not’. Max did *not* say what Pia said him to, given the understanding of being a grunter on which she will have spoken.

The second test is trivially passed. Max said truly, ‘Sid grunts’. Zoë said truly, ‘He (Sid) does not’. There is no occasion on which one can say truly, ‘Sid grunts and he does not’, thereby telling both the truth that Max did and the one that Zoë did. (The qualification rules out the case where one says ‘He does and he doesn’t’, meaning, roughly, ‘Sort of’.)

The third test is also passed, though at the cost of some tedium. Max could quite correctly say, ‘Sid grunts. Of course, one would describe him falsely in those terms (in saying of him “He grunts”) if (as might be) what one were thereby saying is that he is in that habit.’ Not that I advise so speaking to non-philosophers.

#### 4.2 Blocking Communication

What communication requires, C&L suppose, is that from the fact that Max said, on some occasion, ‘Sid grunts’, of Sid (at time *t*), those words meaning in Max’s mouth what they do mean, one can extract what it is that Max thus said to be so. RC of course denies that one can do that. But why in the world should one need to? Perhaps their idea is this. If Max’s words are to be any use to me, I must be able to identify precisely the proposition they expressed. To do that would be to identify precisely *which* understanding they bore; to distinguish that understanding of ‘grunts’ on which they spoke from every other possible and distinguishable one. But (their idea would continue), to do that you would really need to be there (at Max’s speaking). On the RC view, though, communication

just isn't like that. True, Max's words are worth little to me if I know *nothing* of the circumstances in which he spoke. But knowing *something* of them may be enough to know something of what is to be expected if things are as he said. Which may, if things go well, make them quite useful to me. I know he spoke reporting Sid's medical condition. That may tell me all I need to know as to how to understand his words. The idea of identifying 'the' proposition he expressed can simply drop out of the picture.

### 4.3 Contradiction

C&L think RC contradicts itself. It is hard to fall into contradiction without falling into ever so many. But I will concentrate on one. RC denies that sentences have truth conditions. C&L think that RC then goes on to provide some.

To proceed with this charge we will need some idea of the sort of condition RC thinks can be provided. It amounts to this: The English 'He grunts', said of someone at a time, thus meaning what it does mean, says what is true of that person at that time just in case that person is a grunter on that understanding of being one on which one then would speak. There are now two points. The first is that a sentence's truth condition would be a condition on the truth of *that sentence*. That is what RC says sentences do not have. But no condition on the truth of a sentence has yet been stated. What has been stated is a sort of condition on the truth of what would be said in a given speaking of a sentence. I refer back to agreed ground rules. We must not assign a sentence a property it does not have. And *it* does not have a property which some literal speakings of it would have but others would not. The sentence 'He grunts' certainly does not have the property of being true of someone at a time if that person is a grunter on such-and-such non-obligatory understanding of being one.

The second point is a bit more complex. RC's targets—such figures as Donald Davidson and David Lewis—took it that in stating a truth condition for something one was saying when it would be true. One says, '“He grunts” is true iff he grunts', and, the idea is, *that* is when it will be true (namely, when he grunts, supposing that to be some determinate condition of things). That idea could be put in terms of relief from reliance on intuition. Suppose (impossibly) I am a monolingual Finn. I hear Max say 'Sid grunts'. I have not a clue as to when what Max said would be true. Now you tell me what Davidson hoped to say about that sentence: it is true iff 'Sid grunts' (these last scare quotes). Armed with what you told me I need no longer rely on ideas (implicit or not) as to how language works. I have been told how the *world* will be if Max spoke truth. I am in a position to check whether the world *is* as it must be for that. Suppose, now, that, on an occasion, Max said, of Sid, 'He grunts'. Now consult the 'truth condition' that RC supplies for that. When will things be as Max said? You cannot tell. Such information is not extractable from that condition. What you need still to know is how, in the circumstances of Max's speaking, one would

understand talk of being a grunter (what one would then understand being that to be). For that you need your intuitive grasp of how language works. You have not been provided the relief Davidson and followers meant to provide. RC does not contradict itself on this point.

## 5. Minimalism

The driving force of RC is this idea: the open sentences of language speak of ways for things to be which admit of understandings (as evidenced in the second style of argument of section 2). This blocks truth-conditional semantics. For suppose I say, ‘The sentence “Sid grunts” is true iff Sid grunts’. Either I use that last ‘grunts’ on some particular understanding of being a grunter—one understanding among many—or I do not. If I do, then I assign the sentence a property it does not have. For *it* does not speak of being a grunter on any special understanding of this. But if I do not, then I fail to state *any* condition under which anything might be true. Being a grunter on no particular understanding of being one is just not a way for Sid to be. In brief, the choices here are falsehood or failure to say anything. What would be needed to block this result are ways for things to be, which one might speak of, and which do not admit of understandings.

It is central to semantic minimalism that (for given values of referring devices, tacit or otherwise) a sentence expresses a ‘minimal proposition’, where this is something truth-evaluable. (See p. 155 *et passim*.) There is thus a condition on which, and one under which, the minimal proposition would be true. So, too, then, the sentence. If that is what C&L think, one would expect them to have in mind some ways for things to be to speak of which do not admit of understandings. Or at least to have some strategy to lift the barrier RC sets out. So far as I can see, their only strategy here is studied silence. They utter certain forms of words—‘The sentence “Sid grunts” expresses the proposition that Sid grunts, which is true iff Sid grunts’ will do as illustration. But they emphatically refuse any explanation as to what it is they thus said. What they pretend to have said is *what* proposition ‘Sid grunts’ expresses, and *when* it would be true. But if one feels left in the dark by that, they have no help to offer.

C&L anticipate perplexity at this strategy. But they mistake the question that would be raised here. Here is their idea of that question (p. 158):

[The proposition that A is red] is just the proposition *that A is red*. But which proposition is that? What, for example, is it to be just red? What is it that all these things (those that are red when seen through red sunglasses, red on the inside, red in the dark, red on the outside, red when washed, etc.) have in common? What is that state of affairs that the proposition *that A is red* picks out?

. . . this worry has nothing specifically to do with *redness*. It generalizes. What, for example, do all dancers have in common? Some fly in the air, some underwater; some



with music, some without: some stand on their feet, some crawl, some touch other people, some don't. What is the state of affairs that the proposition *that A dances* picks out?

Their question is thus what is in common to all those things each of which is red on some understanding or other of so being. The answer is probably nothing; or nothing *else*. It might be interesting to ask why speaking on all those understandings counts as speaking of being *red*. But that is not the question that arises at this point.

The RC—I, for one—simply wants more details as to when, say, 'the proposition that Pia dances' would be true. If Pia is a trained, skilled, dancer but refuses now to do it, is that proposition true then? If she eagerly takes the floor at every opportunity, but is so clumsy that one might refuse to call it dancing, is it true then? If she dances, but only when you heat the floor enough, is it true then? And so on. The questions are an embarrassment for the reason already stated. If, say, C&L say 'Yes' to the first one, then *that* proposition cannot be the one the English sentence 'Pia dances' expresses. For it does not, as such, speak of dancing on that understanding of engaging in it. If they systematically say 'Don't know', or steadfastly maintain silence, then (unless they are hiding something) there simply is no answer to the question when that would-be proposition would be true; in which case it is no proposition at all.

What is in common to blood in her veins (needing oxygen to redden it) and Pia's lips (reddened by lipstick)? How does that matter to when C&L's 'proposition that A is red' would be true? Not at all. Unless they mean that proposition to be true just in case A is red *on some understanding or other* of something's being so. If they mean some such thing, they are in gross violation of our agreed rules. Suppose we ever did, in saying, 'Pia dances', say what was true just in case Pia dances on never mind what understanding of so engaging. So to speak would be to speak on a *very* unusual understanding of being a dancer. The rule was: do not assign English expressions properties *they* do not have. As we saw, this rule means that if what an expression speaks of being, or doing, is such-and-such, then it does *not* speak of that on such-and-such special understanding of it, so does not say what it would be true to say only on some such understanding. If this is what C&L have in mind, then English simply does not do what they make it out to do. Minimal propositions lead nowhere.

# 7

## Aristotle's Condition

Tim Williamson is the best sort of provocative philosopher. What he provokes is thought. He has made *me* think long and hard about what I am most convinced of. I would like to provoke someone else as much. I would settle for less.

I have a point to make here about truth, and then several others about meaning. For the point about truth I need not invoke my own views on that second topic. It is a point on which I disagree with Williamson. Specifically, I will suggest that bivalence does not hold *as he conceives it*. As for bivalence as one should conceive it, I do not understand what it would be for that not to hold. Bivalence, on Williamson's conception, seems essential to (the motivation for) his own account of vagueness. Perhaps, then, my point threatens that. If so, this is not because of any qualms on my part about there being facts one could not know.

As for meaning, the problem is to locate Williamson's view of that. In fact, he has suggested two views, which move in opposing directions. My own view rules out either of these. But the present project is not to adjudicate. It is rather to see where the crucial issues lie. One of Williamson's suggested views severs the link one finds, for example, in Donald Davidson (e.g., 1967) between what words mean and the conditions under which they would be true (of something). Thus far, Williamson and I are on the same side. The other view has the same shape as a later view of Davidson's (1983), and, for the same reasons, will not (in my view) do. Knockdown arguments to that effect, though, are not here in the offing.

### 1. Aristotle

Williamson, following Tarski, begins a discussion of truth by quoting Aristotle (*Metaphysics* [ 7. 27]):

To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, or of what is not that it is not, is true. (As quoted in Williamson and Andjelković (2000), p. 214, and in Williamson (1994), p. 188)

So someone who says of what is that it is, or of what is not that it is not, speaks truth. And, presumably, conversely (since Aristotle means this as a definition). If

words, or an utterance, or anything else, say(s) of what is that it is, or of what is not that it is not, then they are (it is) true. And presumably conversely. *Mutatis mutandis* for falsity. Something is either true or false (someone speaks either truth or falsity) on condition that it (he) satisfies either the stated condition for truth or the stated condition for falsity. I will call that disjunction of conditions *Aristotle's condition*. It might be expressed this way: to say of what is either that it is or that it is not; or to say of what is not either that it is not or that it is. If the pig is eating turnips, then saying it to be eating them satisfies Aristotle's condition, as does saying it not to be. If the pig is not eating turnips, Aristotle's condition is satisfiable in just the same ways.

One might find substance in Aristotle's condition if one thinks this way. To represent things as a certain way is to impose, or deploy, a particular scheme for categorizing things being as they are: their being that way places them in the one category, their not in the other. It is open to the world to oblige such representing by articulating into things being the way in question, or things not; or, again, to fail so to oblige. If the world obliges, Aristotle's condition is satisfied. If not, not.

So far, I think, Aristotle has an important insight. He continues, 'so that he who says of anything that it is, or that it is not, will say either what is true or what is false'. Williamson endorses the remark. To my mind, the insight just got lost.

There are things to do besides satisfying Aristotle's condition. Everything a drop of water does is something else. So, usually, is lighting a cigar, or going for a walk. Uttering some words need not satisfy that condition either if, for example, they are arbitrarily chosen. Williamson thinks, though, that there is something else which guarantees satisfaction of that condition. It is saying something (to be so). To say that P, on his view, is *always* to satisfy that condition. So it is always either to say what is true or to say what is false. Informally, this is what Williamson means by bivalence. But what is so special about saying that P? What allows it to provide such guarantees? Why is it thus different from, e.g., saying 'P'? I am about to sketch a case that there is nothing thus special about saying that. To do this I will employ 'say that' in one natural way. I think we could, in fact, carve out another way of speaking of saying that on which saying that P *would* entail satisfying Aristotle's condition. That would make it harder (in general) to establish that one *had* said something to be so. And it would make problems erupt for Williamson in a different place. But it is best to stick to just one use of 'say that' for a start. We can turn to alternative formulations of the core point once *some* formulation of it is in hand.

Whatever the reasons for thinking that saying that *is* special in the above way, there are reasons for thinking it is not. For a start, there are cases where, intuitively, it seems not to be. I will sketch a few. They will tend to be far-fetched—of course, since we are prepared for the expected, and (normally) take measures accordingly so that Aristotle's condition will be satisfied.

I. Consider being a bachelor. To win its campaign against premarital sex, the State of Oklahoma passes a law marrying everyone born in the state at birth.

(Some random method is used to pair up—of course—babies of opposite sexes.) These marriages are not recognized in any state west of the Rockies, or east of Dubuque and north of the Ohio, and in few others. Roy was born and raised in Oklahoma. So, in that state, he is married to one Laverne, who he has never seen. Looking to better his prospects, he arrives, at age 30, in Providence, Rhode Island, where he proceeds to cut a fine figure. Is Roy a bachelor? (A pressing question for several of the female persuasion.) Well, yes and no. He is in a way, and he is not in a way. You can't just say that he is, or that he is not. Neither is exactly true. We are accustomed to categorize people as bachelors or not. But, thanks to the State of Oklahoma, the world is not co-operating with our efforts in that direction.

Case Ia. Forget Oklahoma. Roy, a native of Providence, is unqualifiedly single, and eligible in nearly every relevant sense. It is just that there are these little pills. With them, one can change gender overnight. Roy avails himself of these rather indiscriminately. So one day it is Roy, the next day Royine; one never knows quite what to expect. The English 'bachelor' is meant to speak of *males*. So is Roy a bachelor? Again, yes and no. Thanks to science, this time (better things for better living through chemistry), the world refuses to co-operate in this classifying project.

II. Sid has once again lost an expensive leather wallet. For the last time, he swears. Rather than replacing it with another, he constructs 'wallets' out of heavy paper envelopes—a few small ones, for credit cards, business cards, and the like, pasted to the inside of a larger one, where Sid keeps banknotes. Such is what he now uses for a wallet. The new policy, though, does not improve his attention to what he is about. Last night, he left his 'wallet' in a taxi. Overhearing Sid once again engaged in cancelling his credit cards, Pia reports, 'Sid has lost a wallet again.' Is this true or false? Well, he did lose a wallet, and he didn't. He did if we accept his improvisations as wallets. He did not if, on his new policy, he doesn't use a wallet. There is no one right ruling as to how we must describe him in this respect. So it is not quite true to say that he lost a wallet (full stop); but also, equally, not quite true to say that he didn't. To say him to have lost a wallet (at least in the way Pia did) is not unqualifiedly to say things to be as they are, but nor, equally, is it to say things to be otherwise.

III. Sid slams the door shut. Except that he slams it so hard that it shatters into a thousand shards. There is no door at all left in the doorway. Did Sid shut the door? As in the above, he did and he didn't. He did something that normally would count as shutting the door. (The door reached the shut position; otherwise it would not have shattered.) In this case, though, there is reason not to count what he did as shutting the door. Normally, if one shut the door, then (until further notice) it *is* shut. Not this time. So there is good, normally conclusive, reason for counting what he did as shutting the door; but equally compelling reason for refusing to. (The door is not shut.) If that 'equally' is right, then neither policy can be correct as such. So neither that Sid did shut the door, nor that he didn't, can count as true *tout court*.

IV. Finally, a very simple, and very hypothetical, scientific case. Since we are discussing Aristotle, I will set it in a fictitious Aristotelian time. In this time, people (or scientists) thought in terms of an undifferentiated notion of quantity of matter. They thought they had developed two equally good methods of measuring this quantity: spring scales and balance scales. In their environment, the two methods always yielded the same result (within reasonable margins of accuracy). Later, though, it came to be noticed that in certain other environments, for example, at great heights, the two methods diverged. The balance scales continued to deliver the same results. But the spring scales delivered different results at heights than they did at sea level. It thus became apparent (I am supposing) that there are really two different quantities; one measurable by spring scales (and relative to a height) and the other measured by balance scales (and had *tout court*).

The original scientists took themselves to be thinking of a quantity, had by an object in a given degree, full stop, and measurable in two ways. They might, in fact, have been thinking of the one of the above two quantities, or of the other. But there *might* be no fact of the matter as to which they were thinking of. Suppose that an original scientist would have used the words ‘O peseert n livros’ to say object O to have the quantity he had (or took himself to have) in mind to degree n (measured in *livros*). Suppose that the scientist was, determinately, speaking of the quantity measured by balance scales. Then all is well. Suppose he was, determinately, speaking of the quantity measured by spring scales. Nothing has *that* quantity to any degree outright. An object can only have that quantity (in degree n) relative to a height. If the scientist said O to have that property (in degree n) outright, then he did not speak truth. One cannot speak truth in saying any object to have that property outright. But suppose there is no fact of the matter as to which property it is he spoke of. You could take him to have spoken of the first one, or you could take him to have spoken of the second. Then there is no fact of the matter as to whether things are as he said, or not. So he cannot count as having spoken truth. But, equally, he cannot count as having spoken falsehood.

If you could call Roy a bachelor, and you could say he was not, each thing one would say equally consistent with things being as they are, then, so far, neither course could be (as such) saying what is so. For if it were, the opposite course would be saying what is not so. But each, as well as the other, might have either status. Nothing chooses which should have which. So neither course could be saying of what is either that it is, or that it is not. Neither could satisfy Aristotle’s condition. *Mutatis mutandis* for the other cases. So Aristotle’s condition *may* be failed by what says (asserts) something—that P. If, in a given case, it is not failed, then *substantial* thanks are due the world for permitting itself to be spoken of, truly or falsely, in that particular way. Or so the intuition goes.

There is, as noted, another way of speaking of saying that. One could so deploy this notion that in all these cases nothing was said to be so. In that case

it is not decided whether something was said to be so in any given case unless it is decided independently that Aristotle's condition is satisfied. If things here are what they seem, then the world may upset satisfaction of that condition. Nothing is said (on this alternate way of speaking) if the world has done so. So *saying that* will be a highly world-involving affair. This alternate way of talking does nothing to diminish the importance of Aristotle's condition (if these cases are what they seem).

Accepting appearances, Aristotle's condition may be failed even where, on our initial way of speaking, something was said to be so. Pia called Sid a bachelor. So she said that he was one. That is saying something so. As fate has it, though, the condition is failed. And so on. How must we then think of (something) being true? For the moment, suppose truth to be a property of sayables: that is, of something one might say in saying something so. If to be a sayable is *ipso facto* to satisfy Aristotle's condition, then satisfying that condition cannot be part of what separates *true* sayables from others. It is not something demanded of a sayable if that sayable is to be, moreover, true. In that sense it is not *truth* that requires this. Things are otherwise if Aristotle's condition can be failed in the ways just indicated. In that case we can decompose truth, and falsity, into two independent and substantial elements. The first is satisfying Aristotle's condition. On our present notion of saying that, not all sayables do that. The second, in the case of truth, is a condition I will call *merit*; and, in the case of falsity, one I will call *demerit*. Think of merit as not saying things to be other than they are. Think of demerit as saying things to be not only what they are. Failing Aristotle's condition, one may still not have said things to be other than they are; so satisfied the condition on merit. Satisfying Aristotle's condition, what one said may still lack merit. *Mutatis mutandis* for demerit. Writing 'A( $\xi$ )' for *satisfies Aristotle's condition*, M( $\xi$ ) for *satisfies merit*, D( $\xi$ ) for *satisfies demerit*, T( $\xi$ ) for *is true*, and F( $\xi$ ) for *is false*, we now have:

$$(T1): T(P) \leftrightarrow A(P) \& M(P)$$

$$(F1): F(P) \leftrightarrow A(P) \& D(P)$$

So something may fail to be true either because it fails Aristotle's condition, or because it satisfies Aristotle's condition but says things to be other than they are; and it may fail to be false either because it fails Aristotle's condition, or because it satisfies that condition, but says things to be no other than they are. Now writing 'S( $\xi$ ,  $\zeta$ )' for *utterance  $\xi$  says that  $\zeta$* , we have:

$$(T2): T(u) \leftrightarrow S(u, P) \& A(P) \& M(P)$$

$$(F2): F(u) \leftrightarrow S(u, P) \& A(P) \& D(P)$$

So if P does not satisfy Aristotle's condition, then P is not true, and P is not false. Similarly for utterances which say something to be so but fail Aristotle's condition. On our present way of thinking, some values of P may fail that condition. They may be expressed by some utterances which say something to be so.

In assigning such substance to truth, how far do we stray from bivalence? One version of bivalence would be: any  $P$  is either true or false. Schematically,

(B1): For any  $P$ ,  $T(P)$  or  $F(P)$ .

Williamson defines bivalence for utterances. He says, 'If  $u$  says that  $P$ , then either  $u$  is true or  $u$  is false' (Williamson 1994, p. 187) Schematically,

(B2): For any  $u$ ,  $S(u, P) \rightarrow T(u)$  or  $F(u)$ .

If Aristotle's condition has the substance presently envisioned for it, then neither of these versions of bivalence can be correct. But, as we will see, that result puts classical logic in no real jeopardy.

## 2. Classicism

There are cases, I have suggested, where if asked 'P?', we would reply, 'Well, yes and no'. These are cases where we would not be prepared to assert that  $P$ , or, again, that not- $P$ . Neither assertion seems to fit the case. So, equally, we would not be prepared to assert that  $P$  is true, nor that it is false. On the view that I am suggesting, that is because it is not so, in those cases, that  $P$  is true, or that  $P$  is false. Aristotle's condition is not satisfied for  $P$ ; both its truth and its falsity require that. Satisfying that condition is thus a substantive part of what being true is. What stands in the way of taking that idea seriously?

Here, I think, is the main thing. Assume  $P$ . If we have that much, what more could we want for  $T(P)$ ? I confess I cannot think of anything. So  $T(P)$  follows. Now assume  $T(P)$ . Assuming that much, how could it possibly fail to be that  $P$ ? I do not know. If that did fail, then how could it be that  $T(P)$ ? So  $P$  follows. By conditional introduction we thus have  $P \rightarrow T(P)$  and  $T(P) \rightarrow P$ . Combining these by classical logic, we get:

Pseudodisquote:  $P \leftrightarrow T(P)$ .

('Pseudodisquote' because  $P$  is not made up of words; it is a sayable.) Given classical logic, this gives us: not- $T(P) \rightarrow$  not- $P$ . But not- $P$ , like  $P$ , entails  $A(P)$ . (To say that  $P$  would be to say of what is not that it is.) I suggested that for some value of  $P$ , perhaps not- $A(P)$ . But not- $A(P)$  entails not- $T(P)$ , so, we now see,  $A(P)$ . So not- $A(P)$  entails  $A(P)$ , hence a contradiction. So where  $P$  is a sayable, it is contradictory to suppose not- $A(P)$ , just as Williamson suggests. (If some substitution for  $P$  were not a sayable, then, presumably, pseudodisquote would not hold of it, which would block the present untoward result.)

What might block this unwanted result? Tinkering with classical logic *might* do that. In the above reasoning, for example, trouble started with contraposing one side of pseudodisquote. To do that we need negation introduction. And this seems to have a feature we do not want. Suppose that from  $P$  we can

derive a contradiction. Then we had better not assert  $P$ . But there are two ways for  $P$  to be defective. If a statement that  $P$  fails Aristotle's condition, then we had better not assert  $P$  if we aim to say things to be as they are. Again, asserting  $P$  may satisfy Aristotle's condition, but then, not- $P$ . Negation introduction moves us directly to this last case. From the contradiction we may infer not- $P$ . But how did that other sort of failure get ruled out? Perhaps we should not be so quick to conclude what negation introduction says we may.

But such local tinkering seems to lack the required generality. Negation introduction is one route to unwanted results. Are other inference rules more sensitive than it to the possibility that Aristotle's condition is failed? I think not. In our informal argument for pseudodisquote we saw that to assume  $P$  is already, in effect, to assume that  $P$  is true. No gap remains between what we have assumed already and what would need to be so for  $P$  to be true. Similarly, in assuming  $P$ , in some proof, we already assume that Aristotle's condition has not been failed. If it were, it would not be right to assume  $P$ . I cannot coherently suppose that  $P$ , but not  $A(P)$ . And similarly, again, wherever we assume some complex element of which consists in speaking of  $P$  as so (for example, where we assume that if not  $P$ , then  $Q$ ).

Which points to what is needed here. When we are reasoning about Aristotle's condition, or, again, about truth, what needs revising is the rule of assumption (or premiss introduction). We need to insist that it be made explicit, in a proof, just how, or where, satisfaction of Aristotle's condition has been tacitly assumed. So the new rule of assumption should be (at first approximation): wherever there is an assumption, there must be, for each atomic formula in that assumption, the further assumption that that atomic formula satisfies Aristotle's condition. If  $P$  occurs in some assumption, then we must operate under the further assumption,  $A(P)$ . More exactly, and to prevent regresses of required assumptions, the new rule of assumption should insist that a rule of inference can be applied to an assumption only under the further assumptions that that assumption's atomic constituents each satisfy Aristotle's condition.

Has classical logic been revised? Yes and no. (Another example of Aristotle's condition doing work.) One might say that the rule of assumption is integral to classical logic. We have certainly revised that. So, on that way of looking at things, yes. On the other hand, one might look at what we did as simply revising the way classical logic applies to given specific sayables. In fact, perhaps not even that. Perhaps we have merely made explicit something that was always implicit in the correct uses made of classical logic. In any event, in an important respect the revision is not radical. Suppose that  $\phi$  is a classical theorem containing propositional variables  $P_1, \dots, P_n$ . Then our revision leaves us with  $A(P_1, \dots, P_n) \vdash \phi$ . Suppose that  $\lceil \vdash \Delta$  is a classically valid deduction involving atomic formulae  $P_1, \dots, P_n$ . Then  $A(P_1, \dots, P_n), \lceil \vdash \Delta$  remains valid. So, in particular,  $A(P) \vdash P \vee \text{not-}P$ .



As for truth and Aristotle's condition, we have  $A(P) \vdash T(P) \vee F(P)$ ; and  $A(A(P)) \vdash \text{not-}T(P) \rightarrow (A(P) \rightarrow \text{not-}P)$ , as per the substance I have assigned Aristotle's condition. Also,  $P \rightarrow A(P)$  and  $\text{not-}P \rightarrow A(P)$ . We can also now see that my intuitive argument for pseudodisquote should have led to this conclusion:  $A(P) \vdash P \leftrightarrow T(P)$ . None of these things are theorems, or laws, of *logic*. Logic does not unfold the laws of truth in that way. Unfolding the content of Aristotle's condition in mentioning it, as in the above, is the task of a specialized bit of conceptual analysis. Similarly for *so* unfolding the content of truth. (Frege would have insisted: logic does not *mention* truth, or Aristotle's condition, or it could not be the *most general* science. Rather, it *reflects*, in the structure of its laws, some of truth's structure.) Logic is none the worse for such division of labour.

Assuming not- $A(P)$  led to contradiction. But that was by appeal to pseudodisquote. Such appeal now operates under the assumption  $A(P)$ . That *this* leads to contradiction is no surprise. Normally, we apply classical logic where satisfaction of Aristotle's condition is not in question. In such cases we can ignore the modified rule and follow the standard one. Satisfaction of Aristotle's condition goes without saying. Where that condition, or, again, truth, are what we want to speak of, tacitly supposing what the idea of a proof does suppose about them may encourage mistakes as to the source of contradiction.

One might object here that logic is concerned with the most general inferential structure of systems of sayables—a structure that would be there whether these sayables satisfied Aristotle's condition or not. But that structure is just the structure there would be if all relevant sayables satisfied Aristotle's condition. With the revised rule of assumption, logic continues to speak to that. Even if Roy is not quite a bachelor, but equally not quite not one, it *may* remain so that if he were one, he would be unmarried. (Though if I knew that Aristotle's condition were failed by the proposition that Roy is a bachelor, I might well lose some of my faith in that entailment).

When we see something not to satisfy Aristotle's condition, we will no longer hold to it in guiding thought and conduct. We will find other, more adequate, ways of capturing the relevant aspects of things being as they are. Moreover, so long as atomic formulae stand for what speaks of things being thus and so, the first concern of logic will be, not with representations (though, perhaps, facts about representations can be derived from what it says), but, rather, with ways things are and ways things are not—with the sorts of things a successful representation would represent as so. The fact is that the roast is burnt; *that*, and not some representation of it, is what means that we shall dine out. There is something for logic to get a grip on, in the way things are, just where, in representing it as so, Aristotle's condition would be satisfied.

What is at stake here is how we apply logic, not logic itself. In applying classical logic to some given set of statements, we assume things as to their satisfying Aristotle's condition. We have seen how that idea is built into classical logic itself. Where such things as satisfying that condition are not subjects of discussion, we

need no revised premiss introduction rule. In the cases at issue here, what we want to do is apply logic to a discussion of the satisfaction (or not) of Aristotle's condition (or of other successes, such as truth, which entail such satisfaction) by some given set of statements. For those purposes, we (of course) do not want to build it into the very application itself that these statements just do satisfy that condition. What we might take as unproblematic are *statements* that these statements do (or do not) satisfy that condition. In fact, we had better do so if we are to get anywhere. Where (as may be) their satisfying the condition *is* unproblematic, we *may* do so. So far, we have envisioned a case in which there are some statements, P, Q, etc., under discussion. Their satisfaction of Aristotle's condition is part of what is under discussion. But whether that they satisfy Aristotle's condition satisfies Aristotle's condition is not. If we change the case, so that the satisfaction of Aristotle's condition by these further sayables is also under discussion, then they, too, should fall into the class of statements which require our revised premiss introduction rule.

We have now arrived at this position. On the one hand, bivalence as Williamson states it does not hold. That is because he has assigned the wrong significance to saying that P. On the other, classical logic remains intact so far as its theorems and rules for drawing inferences are concerned. What has changed if anything, or at least what is now explicit in the logic itself, is a view as to how logic applies to thoughts (in present terms, to sayables). Logic remains in the same business as always. At most we have revised slightly our conception of what that business is.

### 3. Paradox

It would be rash to claim to have a way of blocking all semantical paradoxes. But many start from a premiss meant to capture a way in which a certain proposition, occurring in that premiss, is about truth, or some related property. To assume that premiss is, as always, to suppose that its atomic parts satisfy Aristotle's condition. So it is interesting to see what happens when that supposition is made explicit according to our revised rule of assumption. (Of course, if there simply is no proposition satisfying the condition that the premiss states, then the relevant argument cannot get started, and there is no paradox. That would be an alternative account of what is happening in any given case.)

I begin with the most familiar paradox. Let *ik* be the proposition that *ik* is not true. So we have:

$$ik \leftrightarrow \text{not-T}(ik).$$

By pseudodisquote, we also have:

$$ik \leftrightarrow T(ik).$$

Which gives us  $T(ik) \leftrightarrow \text{not-}T(ik)$ , from which contradiction follows by the usual route. With our new rules, though, all of this happens under a further assumption,  $A(ik)$ . (And  $A(T(ik))$ , a consequence of  $A(ik)$ .) So we do not have a contradiction on no assumptions, as in the usual presentation of the paradox. So we do not yet have, officially, a paradox.

One move left open here is denying  $A(ik)$ . Applying standard logic, we could derive its negation from the contradiction derived already. Which would cancel that assumption. (Of course, to do that derivation we would need to assume  $A(A(ik))$ .) Now, though, we may seem to be in new trouble. For  $\text{not-}A(ik) \rightarrow \text{not-}T(ik)$ . But, as we have seen already,  $\text{not-}T(ik)$  entails a contradiction. So we arrive at a contradiction on no assumptions. But, on second thoughts, no we don't. To get to that contradiction we need to reinstate  $A(ik)$  (since  $ik$  occurs in both of our initial conflicting premisses). So what we actually derive is a contradiction from  $\text{not-}A(ik)$  on the assumption  $A(ik)$ . As one might write this,  $A(ik) \mid \text{not-}A(ik) \rightarrow (T(ik) \& \text{not-}T(ik))$  (if that is the contradiction we choose to derive). But this is not paradox. It is just an expected result. So we may rest with our conclusion,  $\text{not-}A(ik)$ . (We *could* also refuse that conclusion by denying  $A(A(ik))$ . *Logic* does not force a choice between these tactics.)

A different threat of paradox might appear in the following way. Let  $ik$  be the proposition that  $\text{not-}A(ik)$ . So we have:

$$ik \leftrightarrow \text{not-}A(ik).$$

We also have, by the nature of Aristotle's condition,

$$ik \rightarrow A(ik).$$

From these we can derive  $\text{not-}A(ik) \rightarrow A(ik)$ , from which it follows that  $\text{not-}A(ik)$  entails a contradiction. From which we can derive  $\text{not-not-}A(ik)$ , or, since our inference rules remain classical, simply  $A(ik)$ . From  $A(ik)$  nothing particularly untoward follows. So, it seems, we have proved  $A(ik)$ .

It may well seem, on reflection, that we have not proved very much here. For the derivation from the premisses to  $\text{not-}A(ik) \rightarrow A(ik)$  must proceed on the assumption  $A(ik)$ . So our conclusion,  $A(ik)$ , also rests on that assumption. An argument for  $A(ik)$  proceeding under the assumption that  $A(ik)$  might plausibly be taken to leave the question whether  $A(ik)$  open. So let us ask whether  $A(ik)$ . Nothing may seem to dictate the one answer or the other. But then, since  $ik$  is supposed to say (in some sense)  $\text{not-}A(ik)$ , and since nothing settles the question whether things are that way, that might lead one to conclude, plausibly, that  $ik$  does not satisfy Aristotle's condition, so  $\text{not-}A(ik)$ . I do not present this argument as decisive. Suppose, though, that you feel moved by it. Trouble very obviously brews. For since  $\text{not-}A(ik)$  is what  $ik$  was meant to say, it would seem that  $ik$  does thus satisfy Aristotle's condition after all. So we are stuck with both  $A(ik)$  and  $\text{not-}A(ik)$ . Our plausible conclusion thus seems to entail a contradiction.

If this informal argument were made rigorous, under present constraints on derivations, it would be seen to depend on the assumption  $A(ik)$ . In that form, it cannot quite count, as it stands, as paradox. Still, neither that  $ik$  satisfies Aristotle's condition, nor that it does not, seems a satisfactory thing to say. If I try to say that it does not, I presuppose that what I am saying satisfies Aristotle's condition. But since that seems to be what  $ik$  is saying, it would seem that I am supposing that  $ik$  satisfies that condition too. I do not want to suppose such things. So I must not say that  $ik$  does not satisfy Aristotle's condition. But to say that I must not say it not to is to say that what I would thus say—that  $ik$  does not satisfy Aristotle's condition—does not satisfy Aristotle's condition. Equally, then, for the thought that it does. So we have  $\text{not-}A(A(ik))$ . We might see  $ik$  itself as trying to say something near enough to this to be true if what we say is, namely,  $\text{not-}A(\text{not-}A(ik))$ . But we need not agree that it succeeds in saying this. Not so far, at least. So we might plausibly see this last suggestion as a result.

#### 4. Semantics

So far I have expressed no view about meaning. I have cited intuitions in favour of the role I have assigned Aristotle's condition. That role may rule out *some* views of (linguistic) meaning. For example, it would rule out a view on which the meaning of an open English sentence, in each case, is, or fixes, a function which assigns every object a truth value. But I have not derived my view of Aristotle's condition from any specific characterization of what meaning does, or of the properties an expression has in, or by, meaning what it does. I do have a view on that matter, though, which, I think there would be no room for if Aristotle's condition did not work as I said. As noted, Williamson sketches two different views of meaning, each of which is at odds with mine, though each for a different reason. The rest of this essay aims to identify these points of conflict, and thus the issues on which adjudication rests. It is not meant as a refutation of either of Williamson's views, though if all goes well it may have some persuasive power. I begin with my own view in five points.

I. The English 'is blue' speaks of (an item's) being (coloured) blue. The English 'wrote a novel' speaks of (someone's) having written a novel. The English 'grunts' speaks of (a creature's) being a grunter. The English 'weighs one kilo' speaks of (an item's) weighing one kilo. And so on. These are the sorts of properties a language's expressions, or at least its open sentences, have in and by meaning what they do. At the very least, they are properties such expressions *have*.

II. English expressions (or open sentences) thus speak of what admits of understandings. Being blue, for example, admits of understandings. It is a particular way for an item to be (coloured). But take Lac Lemman on a sunny day. Is *it* blue? It is if you understand (a lake's) being blue in one way, but not if you understand this in another. On one understanding, the reflected sky makes

it blue. On another, since the sky does nothing to change the composition of the water (it looks the same when bottled), the lake is not blue. These are both understandings being blue admits of; things it would not be wrong (*tout court*) to take a lake's being blue to be. Similarly for having written a novel. Our modern authors try to stretch the boundaries, but, in my view, sometimes cross them. Must a novel have narrative structure? Must it be a connected story, or even a story at all? Or may it be (*à la* Pynchon) rambling ruminations, akin to the bar bore's 'And another thing I don't like'? Are de Beauvoir's memoirs, names changed, a few fantasy scenes intentionally added, fictional enough to be a novel? How must novels be distinguished from history or travelogue? As to the writing, may that be ghosted, and to what extent? Or, if Burroughs writes a text, cuts the pages in half, shuffles them, pastes them back together into pages, and we all agree that the result *is* a novel, did he *write* it? There is an understanding of writing a novel on which this is writing one, and an understanding on which it is not. Some of the problems here have to do with borderlines. (How much ghostwriting?) But some, like the last, do not. Similarly for the other examples. And similarly throughout the things spoken of by a language's open sentences.

III. Where an open sentence speaks of (precisely) A, and A admits of understandings, the open sentence, as such, speaks of A on none of these understandings in particular. For if it did, what it would speak of, in meaning what it does, would be, not A, but rather A on such-and-such special understanding of it; which, by hypothesis, it does not. So, for example, The English 'is blue' does not, in or by meaning what it does, speak of a lake's being blue on an understanding on which reflected sky would make a lake that; nor on one on which reflected sky, so far as that went, would fail to do the job. It speaks of being blue in a way that is neutral between these options. Otherwise it would not, full stop, speak of being blue. Which it does. Corollary: What 'is blue' means does not yet determine whether it is true of Lac Lemman on a sunny day, or when it would be true of Lac Lemman, or of anything else. Meaning does not connect with truth like that. Meaning might supply some materials for truth; but some materials must come from elsewhere.

IV. If, on an occasion, you use 'is blue' to say something of a lake, using it to mean what it does mean in English, you will thus have called it blue, or described it as blue, or said it to be blue. If you use it of a lake non-assertively (e.g., 'If Lac Lemman is blue, I shall take fish for lunch'), then you speak of the lake as being blue. That is, if, on the occasion, you want to do what it would then be to call a lake blue, or speak of it as blue, and if speaking English will do the trick, then, more specifically, 'The lake is blue' will do the trick. In the circumstances of your so speaking, there *may* (though need not) be something that would, naturally, reasonably, be understood by a lake's being blue. Sid and Pia, walking along the shores of Lac Lemman, exult at their escape from grey North Sea skies. Sid exclaims to Pia, 'How blue the lake is!' It would surely be churlish for Pia to scoop a cupful of water out of Lac Lemman, show it to Sid, and say, 'No it's not'.

Churlish is an understatement. If she is not joking, then she failed to grasp what it is that Sid said. Insofar as there is such a thing as an understanding there would be of what one would be speaking of in speaking, on an occasion, of (a lake's) being blue, in speaking of its being blue, on that occasion, you speak of its being blue on that understanding. (That is, whether things are as you said turns on whether the lake is blue on that understanding.) So for any understanding (a lake's) being blue bears, if you find the right occasion, you may use the words 'is blue' to mean what they do and thereby speak of a lake's being blue on that understanding. Meaning constrains what is said no more than that. I hope it is clear how the example generalizes to other cases.

V. Therefore, where there is an item that may be sensibly, understandably, called blue, it is, in general, both possible to speak truth, and possible to speak falsehood, of that item, at a given time, in using 'is blue' of it, and as meaning what 'is blue' does. There are truths to be expressed, but also falsehoods to be expressed, in using 'Lac Leman is blue' to say something about Lac Leman-at-a-moment. In any event, what 'is blue' means does not determine when it would be true of what. Again the example generalizes.

That, in brief, is how I see meaning. I note a corollary about the notion of saying that. On the above view, the words 'Lac Leman is blue' are for doing a certain sort of thing. An apt title for that sort of thing is: *saying that Lac Leman is blue*. If you use those words, on an occasion, to speak of Lac Leman, and you say something (to be said in speaking English), then you produce an instance of that thing there is to do. So what you do merits the title *saying that Lac Leman is blue*. This fixes *one* reading of the expression 'say that Lac Leman is blue'. To credit you with doing that, on this reading, is not to credit you with speaking on any particular understanding of a lake's being blue. So it is not to say what fixes when what you thus said would be true. But 'say that Lac Leman is blue' also admits of another reading. This is a reading those words are likely to bear if I tell Pia, 'Lac Leman is dull grey today', and she protests, disappointed, 'But Sid said that it is blue'. I speak on a particular understanding of a lake's being dull grey. Pia is to be understood as speaking on an understanding on which what she credits Sid with saying is inconsistent with what I thus did. To credit Sid with such a thing is to credit him with speaking on a particular understanding of a lake's being blue. Context here provides that understanding. It is one that matches up, in the intended way, with the understanding my words 'is dull grey' bore. So there is also this reading of 'say that'. On it, to say that a lake is blue *is* to say this on a particular understanding. The occasion of that use of 'say that' must determine what understanding this is.

Using 'is blue' as meaning what it does, speaking English, speaking of a given object, is not doing enough to ensure satisfying Aristotle's condition. In fact, if one is speaking of a typical lake, and one does no more than this, one will not satisfy that condition, since a typical lake is blue on some understandings of so being, not on others. So what the words thus used *mean* does not do enough to

make either for saying things to be as they are, or for saying things to be as they are not. One could do that only in speaking on some particular understanding of an object's being blue, thus an understanding which went beyond what is fixed by what the English words 'is blue' mean, or speak of, as such. Perhaps using 'is blue' of an object, speaking English, so that 'is blue' so spoken meant what it does mean, is enough to make for saying that that object is blue, on our first reading of saying that. If so, then saying that P, on this reading of 'saying that', is not enough to ensure satisfying Aristotle's condition. What of the second reading? Suppose you say the object to be blue, speaking on a particular understanding of an object's being blue. If there is not enough in what words *mean* to guarantee satisfying the condition, then neither is there enough in this more specific accomplishment to *guarantee* that. For it may turn out that there are understandings of being blue on the understanding on which you spoke of this, on some of which the object would be blue, on others not, and that nothing in what you did makes it so that you spoke on an understanding of the first sort rather than one of the second sort, or vice versa. Which would place *your* words 'It's blue' in exactly the same position as the English words 'It's blue'. For the same reason that *they* neither say a given object to be as it is, nor say it to be as it is not, as fate would have it, neither do your words, given the understanding they in fact bore.

In the event, Williamson thinks my view of meaning rests on a mistake (or several). He says:

According to Travis, meaning and reference underdetermine what is said . . . However, his account of what underdetermines what is said seems to conflate use and mention. He assumes that someone who says 'That is round' of a ball, using the words with their usual meanings, describes the ball as round, and argues that, given the way the ball is, whether it is true to describe it as round depends on the context of utterance. This commits Travis to rejecting the plausible principle that it is true to describe the ball as round if and only if it is round. But once we accept that 'round' is context-dependent, we should reject Travis' assumption, just as we should reject the assumption that if you say 'That is mine' of something, you describe it as mine. . . . 'Mine', with its usual meaning, does not always refer to the property of being mine; why should 'round', with its usual meaning always refer to the property of being round? With this correction, Travis' argument can be reconstructed as leading to the less radical conclusion that meaning alone ubiquitously underdetermines what is said. However, he simply takes the sameness in meaning of the words in his examples as obvious, and does not discuss the possibility of slight meaning changes, as postulated below. (Williamson 1998, p. 10 n.)

There are here two different views of meaning. I will consider each in turn.

The first view combines two thoughts. First, the English predicate 'is round' does not always speak of being round. (That is, not always of the property I just mentioned in those words, assuming that I did just mention one.) Second, it speaks of (refers to) different properties on different occasions of its use—just as 'mine' speaks of different people on different occasions of its use. I take it

that what goes for 'is round' here goes for 'is blue', and so on, (more or less) systematically throughout English.

First, then, Williamson speaks of 'round' as referring (or, in the event, not) to 'the property of being round'. He thus presupposes that *he* referred to some property or other in *his* speaking of those words. So, sometimes, in speaking English, we do, in his view, refer to, or speak of, properties. Whatever property he did speak of on this occasion, he suggests that the English word 'round' (or predicate 'is round') does not 'always' refer to that one. Well, if it does not *always* refer to that property, then *it* (that bit of English) does not refer to, or speak of, that property, full stop. We can delete the 'always'. The property of speaking of that property is not one that *it* has. If I have the property of being fat, I do not merely have it when speaking to you, but not when speaking to Jones. (Bizarre images come to mind when one tries to imagine that. They are not ones of my being fat.) Nor do I have it merely on some uses of me (citing me in the *TLS*, say), but not on others (citing me in the *Crimson*). Having the property of being fat, when that feat is accomplished, calls for no relativization. Similarly for the property of speaking of being round. Moreover, in the only aspect of 'speak' on which it makes sense to think of a bit of English as speaking of something (roughly habitual, or functional, as in 'Persil gets clothes really white'), for the English word 'round' to speak of some property would be for that to be what you speak of in using it, provided you are speaking English, using its words to mean what they do. 'Round' speaks of that property, as Persil does its thing, *when used* (properly).

So the English 'round' does not speak of the property Williamson mentioned in his words 'the property of being round'. If I am right about properties, or ways for things to be, then if it does not do this for the reasons Williamson has in mind here, it does not speak of any property (or, in my terms, way for a thing to be). For if I am right, just mention a property (as Williamson did), and, as soon as we are clear on what property that is, I will show you how it admits of understandings. In which case, if the English 'round' spoke of that one, then, in so using it, you would sometimes speak of *that* property on *this* understanding of it, sometimes on *that* one. We would be driven back to Williamson's suggestion that 'round' does not (always, so ever) speak of the property you mentioned.

I think that the English 'is round' speaks, as such, of being round. So I think it speaks of that on every use of it which is (a case of) speaking proper English. I am not dissuaded from this view by the fact that, on different such speakings of it, it will make different contributions to the truth conditions of wholes of which it may, then, be part. There is a fact about that bit of English which I hope I can capture in that way. It is that the meanings of those words constrain, in a particular way, what you can say on an occasion in using them—even if they do not narrow things down to just one thing. I can, occasion permitting, call the squash ball round and speak truth of it even as it begins its rebound off the wall. But I cannot so easily call it round and speak truth of it just in case it barks, or



is made of lead, or is on fire. My idea is: I can use 'is round', on an occasion, of a ball, to say what I would say on that occasion in *calling it round*; and (*ceteris paribus*, perhaps) that is all I can use those words to say of it. Further, to say what I just did is to say how the meanings of those words constrain their use. I cannot use those words of the ball (speaking English, so that they mean what they do) and thereby say the ball to be on fire, unless there is an understanding of being round on which to be round is to be on fire. (Which, so far as I can see, there is not.) I do think this is a pretty stringent constraint on what you can use 'is round' to say in speaking English. If Williamson does not want to capture it as I do, if he does not want to say that 'is round' speaks of (refers to) being round, then I am curious as to just how he would capture it.

I turn now to the second component in Williamson's first view. Williamson suggests that 'is round', in, say, 'The ball is round', refers to different properties on different uses of it. So, on different such uses, it makes different contributions to the truth conditions of the resultant whole. In one case, that whole is true just in case the ball has one property; in another the whole is true just in case the ball has a certain different property. And so on. This is to say that, as to the English 'is round', there is no one contribution to truth conditions which is *the* one which it makes. So, in constructing a semantic theory, we can seize on no such supposed contribution in order to identify what it is that that expression means.

Suppose we aimed for a semantic theory in Tarski style, as Davidson conceives this. (See Davidson 1967.) Then we would want to assign a certain bit of English, 'is round', a 'satisfaction condition'. This amounts to adopting (or trying to adopt) an 'axiom' of our theory which would be written like this: 'The English predicate "is round" is true of something just in case that thing is round.' (I adjust terminology slightly here.) The words of the axiom, to the right of the 'just in case', purport to identify those conditions under whose obtaining the predicate would be true of a given thing. More perspicuously (I think), they purport to identify that condition (*Zustand*) of a thing in which it would be such that the predicate was true of it. Varying terminology once again, it purports to identify that property the having of which would make the predicate 'is round' true of a thing. Williamson's idea, if correct, provides a very perspicuous way of saying why this project is doomed to failure. Suppose those words to the right do mention (speak of) a property. That will be one the words 'is round' do not refer to as such. As Williamson puts it, they do not always refer to it. But then the proposed axiom states a generalization which is false. Now suppose that those words to the right do not mention *any* property. Then no condition for the truth of that predicate of something has been stated at all.

If things work as Williamson suggests, then, so far as bits of English are concerned, there is no satisfaction to be had. Thus far Williamson and I are entirely on the same side. Is there anywhere we differ? Williamson speaks of different speakings of 'is round' as referring to different properties. I speak of them as speaking of being round on different understandings of being round.

If there is any difference here, it lies in our conceptions of, respectively, an understanding and a property. On my view, to speak of being round on a particular understanding is to speak such that there are particular things to be understood as to when something would count as being round, where these things are not built into the notion of being round itself. That they are not is visible in the fact that they need not always be understood wherever one speaks of being round. For example, it may be to be understood that the squash ball's momentary geometry, as it rebounds, is irrelevant to whether it is as said *here* in calling it round. In a happy case enough is to be understood of this for it to be clear, given the way the ball is, whether one ought to call what was said true, or false. In an unhappy case this will not be so. (The point of Aristotle's condition.) But for there to be enough to decide what needs deciding in the case at hand need not be for there to be enough to decide everything that ever might need deciding.

I suspect that Williamson builds a bit more into his notion of a property. In fact, he must do so if he is to deny Aristotle's condition substance. If we think of our (shared) environment as inhabited by some determinate collection of objects, and we ignore issues about categories, should there be such, then, on Williamson's conception, a property, as such, partitions this set exhaustively into two disjoint ones: the set of all those objects which have the property, and the set of all those which lack it. In fact, a property (if it really is that) is guaranteed to do this, no matter how things happen to be. This conception of a property underwrites Williamson's conception of bivalence. A statement is true if the objects it speaks of have the properties it ascribes them, and false if some of them lack some of those properties. It is in the nature of a property, so conceived, always to make either the one thing or the other the case. My own view, though, is that this conception of a property is inconsistent with Williamson's supposition that we sometimes speak of them. I do not believe that we know how to identify anything for us to speak of (any way for a thing to be, that is) which does not admit of competing understandings (in ways we are entirely equipped to recognize). I have given here no knockdown argument for that. This, though, is the key point on which Williamson and I differ about meaning.

Williamson's second view of meaning is contained in the last sentence of the above quote. It is that the meanings of words change, sometimes subtly, over time. If there is such a thing as English, and things *its* words mean, then time can be the only relevant variable here. If we want to think of different speakers at a time (both speaking English, if we speak loosely enough about doing so) speaking slightly different idiolects in which words have slightly different meanings, then we might try to relativize meaning to that too. I am sure that words do change meaning over time, and that that is why, for example, the Dutch 'typies' and Portuguese 'esquisito' do not mean, respectively, *typical* and *exquisite*. But I think the meaning changes Williamson has in mind here are

meant to support a certain idea about what meaning does. In the case of a predicate, such as 'is round', that idea would be that, for a *fixed* meaning (its meaning at a moment), there is a unique, invariant contribution it would make to the truth condition of any whole it was a part of while meaning that. Though the meaning of 'is round' may change over time, freeze it at a time, and it does have a satisfaction condition of just the sort Davidson envisaged. *Mutatis mutandis* for relativization to idiolects. As meaning changes, so do conditions on satisfaction. This view *preserves* Davidson's conception of what the meaning of an open sentence does—that it fixes an effective condition on satisfaction—whereas, by contrast, Williamson's first view of meaning abandons that connection.

But for currently irrelevant details as to how an idiolect is fixed, and the importance of speaker intentions (or understandings) in fixing what was, in fact, said on an occasion in given words, this idea about meaning is essentially the later view of Donald Davidson on that topic. (See Davidson 1983.) It blocks the following kind of argument: 'Jones said something true of that ball in saying of it, "The ball is round"; Smith said something false of that same ball (in that same condition) in saying of it, "The ball is round"; hence what "is round" means is compatible both with saying something true, and with saying something false, of a ball in that condition.' For now, all the premisses give us for sure is that 'is round', when Jones used it, meant what was compatible with saying one thing; and, when Smith used it, meant what was compatible with saying another. Which is not yet the result we wanted. For we do not yet have that it meant the *same* on both uses.

Still, the idea here fails to do the work required. For the mainspring of the opposing view was this. Choose any way for a thing to be. Specify it any way that you are able. Then we can recognize various competing understandings of being that way—various competing things it might be for a thing to be that way. That is not, as it stands, a thesis about words. So freeze an idiolect at a time. Pick a predicate in it. Now tell me what that predicate speaks of in *that* idiolect at *that* time. Then I will show you various understandings of something being that way. Which says something about what predicates are for: *not* imposing some given condition on satisfaction, but rather for speaking, on whatever occasion, of whatever it is that might count on that occasion as being thus and so.

One can put this point in terms of something we all share as thinkers, across any variation in our idiolects. My understanding 'is round' as I do, in my perhaps highly idiosyncratic idiolect of the moment, makes me prepared to use it in any of indefinitely different ways, as need arises. My understanding of it (in my idiolect) shows me that it *has* all those different ways of being used, given what it speaks of there and then. In assigning it the meaning I do, I am prepared to recognize all those different things as ways it, so meaning, may be used. What matters here is not the specifics of any particular idiolect, but rather what it is, in general, to say something to be a given way. I do not know the fine details of your idiolect. But I am sure you share with me this general conception of

how words work—if not explicitly, then at least in what you are prepared to recognize. If I only *think* I have the understanding of predication that I do (and not just of this or that predicate), what I need is some considerations showing (or making probable) that this is so. The present idea, common to Williamson and Davidson, even if accepted, would not militate in that direction.

## 5. Boundaries

Frege writes:

A definition of a concept . . . must unambiguously determine, as regards any object, whether or not it falls under the concept . . . Thus there must not be any objects as regards which the definition leaves in doubt whether it falls under the concept; though for us men, with our defective knowledge, the question may not be decidable. We may express this metaphorically as follows: the concept must have a sharp boundary. . . . A concept that is not sharply defined is wrongly termed a concept. Such quasi-conceptual constructions cannot be recognized as concepts by logic; it is impossible to lay down precise laws for them. The law of excluded middle is really just another form of the requirement that the concept should have a sharp boundary. . . . Would the sentence 'Any square root of 9 is odd' have a comprehensible sense at all if *square root of 9* were not a concept with a sharp boundary? Has the question 'Are we still Christians?' really got a sense if it is indeterminate whom the predicate 'Christian' can truly be asserted of, and who must be refused it? (Frege 1903, §56)

Logic, Frege tells us, does not describe the behaviour of concepts without precise boundaries. It would not apply to such concepts mechanically. It would apply to such 'quasi-concepts' only insofar as they are harmlessly idealized in taking them for (Fregean) concepts. It takes a sensitive eye to see just when this would be so. (Compare the way mechanics applies to actual bodies.) He may further *suggest* that without this precision words have no definite sense, which is to say, express no determinate thought, say nothing in particular to be so. I endorse Frege's claim about logic, read as I just have. If you want to apply laws of what has precise (enough) boundaries to what does, or may, not, then you must be careful to mark, or note, the discrepancies between what those laws are designed to hold of, and what you are, in fact, applying them to. Explicitly assuming satisfaction of Aristotle's condition where its satisfaction is in question in an argument is one way to take such care. I dissent from the further suggestion, if there is one. Saying something to be so (saying that P) does not *per se* require deploying concepts with sharp boundaries in Frege's sense. In discussing semantics I have suggested that our concepts do not in fact have such sharp boundaries, so that it is open to the world to decide, in being as it is, whether a given deployment of them satisfied Aristotle's condition or not. Further, I have argued, in the first part of my discussion, that nothing in logic should dissuade us from such a view. As Frege emphasizes, logic is not aimed at *such* questions at all.

Williamson's account of vagueness requires, not just Frege's view of what laws of logic are about, but Frege's further suggestion as to the connection between precision and making sense: that wherever something is said to be so, there will always, of necessity, be a sharp boundary between that which is things being as thus said to be and that which is not. So, where something is said to be thus and so, there will always be a sharp boundary between that which is as that thing was said to be, and that which is not. It is thus this suggestion that divides us.

The suggestion applies to vagueness through cases like these. We arrange paint chips in a series so that they are pairwise indistinguishable by sight in colour, but so that the first member of the series is clearly red, the last clearly orange. Now pick a chip somewhere in the middle. Then Williamson's idea is: of necessity, either it is red, or it is not. That is so even if, as will be if the chip is well-chosen, we do not, and cannot, know which.

It has struck many as absurd to suppose that a chip is red, or, equally, that it is not, if no matter how much else we knew about it, no matter how full our access to its being as it is, nothing available to us would so much as give us reason to think it red, or, as the case may be, not. The operative principle here would be: if thought that things are such-and-such is to be a *truth-evaluable* stance towards the world—a stance suitably correct, or not, by virtue of things being as they are—then the world's being as it is must be able to bear on the question what we are to think *in re* things being such-and-such—whether that is the thing for us to think, or not. So that the chip is red is not truth-evaluable, so not a genuine thought at all, if the world *could* not bear on the question whether we are to think the chip that way. Against this line of thought Williamson has argued, if I understand him, that an epistemologically sound economy of thought would have to make for truths we could not know. I will not approach that argument here. But two brief remarks.

First, we are now in a position to see that a case against the idea that the chip is red, or not, but unknowably so, need not begin from premisses about our ignorance. There are things we *do* know about the sort of chip at issue here. We know that there are truths to tell in saying it to be red, and also truths to tell in saying it not to be. For we may say it to be red on any of many understandings of its being so. Our apparent helplessness in the face of the philosopher's chosen chip is (thus) an artefact of the circumstances of the choosing (namely, doing philosophy). In the circumstances in which we are asked to say whether the chip is red, there is insufficient (or no) point in saying the one thing or the other; no consequences of our so saying, no expectations thus aroused, which might arrange for our speaking (in, say, calling the chip red) either on an understanding on which one of those truths would be told, or on one on which one of those falsehoods would.

Suppose that you say O to be F. Suppose that whether O is F depends, substantively, on what one understands being F to be: there *is* an understanding of being F on which O is F, but also one on which O is not. Suppose there is

no more reason to understand you to have said O to be F on the one sort of understanding than there is to understand you to have said this on the other sort. So neither way of understanding what you said is better than the other. So neither is *right*. Then there is nothing in (or about) what you said that determines, at a place where determination happens to be needed, how the world should be to make what you said true. In such a circumstance the world is powerless to be, in being as it is, either as you said it is, or not. Such is the nature of the gaps that may arise between saying something to be so and, satisfying Aristotle's condition, thereby saying something true or false. When it comes to answering the philosopher's question what colour the chip is, we are discernibly, determinately, in that position.

Second, *perhaps* the idea of saying, or thinking, something either true or false does come apart, in the odd case, at least, from the idea of the world, in being as it is, being capable of bearing on the question whether we are to think that thing—that is, of the world being capable of providing us with reasons we could see to bear on the question whether that thing is so; reasons which actually militated in favour of its being so, or (as the case may be) not. The world's being as it discoverably is may show that there are cases where these two ideas cannot be held together. There is room for discovery (by physics, say) that truth-evaluability is not what one would, at first, have thought it. I am sceptical as to philosophy, or logic, one day showing such a thing. In the present instance, logic (and philosophy) are meant to have paved the way to the parting of these two ideas in this way: logic shows that there is no room for failure of bivalence in Williamson's sense—for something to be said to be so which, for all that fails Aristotle's condition; to call the paint chip red is to say something to be so (it would be poor philosophical methodology to suppose otherwise); hence, to call the paint chip red is to say something either true or false. But, I hope to have shown, logic shows no such thing as the first step in this line of thought.

*This page intentionally left blank*

Part II

Applications



*This page intentionally left blank*

## 8

# Are Belief Ascriptions Opaque?

'A bad question for a title', one might think, 'for words like "believes that John is married"' have two senses; used in one sense they are opaque, used in the other not.' But this is an instance of a ploy which, while perhaps sometimes correct, is generally suspect: where there are two competing, initially plausible analyses of what is expressed in given words, posit two notions, each fitting one analysis, and each expressed sometimes in those words. Experience has shown that it is better first to try to arbitrate between the reasons apparently in favour of each. Take 'red', for example. It seems plausible to take this to indicate a state of things, where science might show which state this is. But it also seems plausible to take it to indicate something to be seen by looking, where nothing else could controvert what looking shows. Since the intuitions have seemed to conflict, philosophers *have* posited two senses of red—something like a 'dispositional' and an 'occurrent' sense, as with C. D. Broad, for example, or perhaps an 'observational' and a 'theoretical' sense. But if we see why each of these intuitions would arise for just one notion, and why that notion should be taken to be that of being red—in this case, something shown, I think, by Hilary Putnam (1975/1962; 1975/1963) and J. L. Austin (1962*b*, esp. lecture X)—we may also see why, in this arena, the best account of the matter does not posit ambiguities. Perhaps there is the same thing to be seen about belief. So, holding questions of ambiguity open, I propose here to review some of what we know about belief ascriptions. I do not expect to settle whether there are opaque ones, since I do not think that we yet know what we would need to know to decide that. I do hope to show that the usual reasons for finding (some) such ascriptions opaque are bad ones, as are the usual accounts of what such opacity would come to. In fact, there is more than just a possibility that in this area, our notions of opacity such as they are, break down.

### I

*Opacity.* Opacity, for reference to particulars, is thought to come in two forms, *substitutional* and *existential*. From now on, I will consider it only with regard to (ordinary) names such as 'Frege' or 'Arthur Dent'. *Substitutional* opacity has something to do with preserving truth while replacing, in some words in which

it occurs, one name by another. A first attempt at a description might look like this: Let XAY be words expressing a truth in which A is a name which refers to O. Let B be a name and XBY such that in it B refers to O. If the truth of XBY is guaranteed by that of XAY, then the *context* X \_\_\_\_ Y (the words XAY minus A, and understood as requiring insertion of a referring device for expressing a thought) is substitutionally (s-) transparent. If there is no such guarantee, then the concept is s-opaque. *Existential* opacity has something to do with removing A from XAY and forming an existential generalization from the remainder. There may be many ways of doing this in English. What such a generalization ought to say, roughly, is that there is something such that what X \_\_\_\_ Y says of something is true of it. Such might be said, e.g., in the words, 'There is something such that' X 'it' Y. If the truth of such generalizations is guaranteed by that of XAY, then the context X \_\_\_\_ Y is e-transparent. Otherwise it is e-opaque. These explanations are crude in ways that will matter. But let the refinements come as the need for them is shown.

## II

*Quine's argument.*<sup>1</sup> Quine holds that s-opacity entails e-opacity. If so, then if we know that a context is e-transparent (i.e., supports existential generalizations), then we can know that it is s-transparent. Which may mean that e-transparency is more difficult to recognize than it looked to be.

Here is what I take the argument to be. Consider a generalization of form 1:

1. There is something such that X it Y.

What would make it true? A plausible answer is: there being an item, O, such that what the words X \_\_\_\_ Y say/said to be so of something is so of O. Now suppose there is an O with two names, A and B, such that XAY is true, but XBY false (thus, X \_\_\_\_ Y is s-opaque). Suppose that what X \_\_\_\_ Y says to be so of something is so of O. Then, it would seem, XBY must be true; counter to assumption. So rule out that supposition. Now suppose that what X \_\_\_\_ Y says to be so of something is not so of O. That would make XAY false, counter to assumption. So scratch that supposition as well. Suppose that what X \_\_\_\_ Y says to be so of something is neither so nor not of O (there is no fact of the matter). Then both XAY and XBY must be neither true nor false, counter to assumption. So that supposition is also out. Since this exhausts the options, we had better deny a presupposition: There must be no such thing as 'what X \_\_\_\_ Y says to be so of something', if X \_\_\_\_ Y is s-opaque. So the conditions for (1) being true cannot so much as be stated, much less fulfilled. What comes to the same thing: if

<sup>1</sup> A good place to look for a statement is Quine 1961/1953. See also Quine 1966/1956.

(1) has well-defined truth conditions (and accepting the plausible account of what they would be),  $X \text{ — } Y$  cannot be s-opaque.

Quine's thesis gains support from examples like his own (2):

2.  $\text{—}$  is so-called because of his size.

Fill the blank with 'Giorgione' and you get a truth. Fill it with 'Barbarelli' and you get a falsehood. A plausible explanation: there is no such thing as 'being so-called because of one's size', hence no such thing as 'there being someone who is so-called because of his size'.

Note the strength of Quine's conclusion. The problem is not that there may be a few strange items, such as  $O$ , about which there are truths such as  $XAY$  which do not support existential generalization. If there is nothing which  $X \text{ — } Y$ , on its own, says to be so of something, then there are *no* true generalizations of form (1), given the plausible account of their truth conditions. Note also what does *not* follow from this. The argument does *not* show that if  $X \text{ — } Y$  is s-opaque, then there may be truths of the form  $XAY$  where  $A$  refers to nothing. From the story about (2), for example, we ought not to conclude that I might express a truth in saying 'Barney is so-called because of his size'—say, because 'Barney' is a subtle allusion to infantile dimensions—even though 'Barney' on my lips refers to no one. Such is not the force of Quine's thesis.

### III

*Pseudo-opacity.* Consider the following situation: There is an extremely dignified business woman named Mary MacAdam, and two nearly distinct groups of people who speak about her. We are unique in belonging to both. One group knew her only as a child, when she was always called 'Mimi', kept pets, and never ate anything cute (she was a near vegetarian). They tell stories, now and then, about her childhood; know nothing of what happened to her since. The other group knows her only as a pet-abhorring, carnivorous adult, called either 'Mary' or 'Ms MacAdam', full stop. Now consider the context (3):

3.  $\text{—}$  had a little lamb.

Fill the blank with 'Mary', and you may find instances produced by the second group. Such could only be understood as describing what Ms MacAdam ate for dinner. So understood, suppose, they express truths. Fill the blank with 'Mimi', and we may find instances produced by the first group. Such could only reasonably be understood as speaking of MacAdam's childhood pet keeping. So understood, suppose, they are false. (She had many pets, and memories dim.) So, by searching transcripts, we can find instances of (3) filled in by 'Mary' which are true, and instances of (3) filled in by a co-referential 'Mimi' which are false. Would this show (3) s-opaque?

The answer, surely, is ‘no’. It had better be if there is to be any transparency left. The trouble, clearly, is that from one filled-in (3) to the other, the contribution of (3) itself to what is said changes—in this case, because ‘had a little lamb’ is ambiguous. If that *is* a trouble, then we may draw the following moral: for words XAY and XBY to exhibit the s-opacity of X \_\_\_\_ Y, X \_\_\_\_ Y must make the same contribution to what is said in each case. Call this the *principle of constancy*. The thought may also be put this way. Where opacity is at issue, it is of no interest to be told that the words X \_\_\_\_ Y may say something which is true of an item O, and also may say something *else* which is false of that item. *So* much is true even of transparent (3). But where the principle of constancy is not respected, mere contrast in truth value of some XAY and XBY, for co-referential A and B, *by itself* can show no more than this.

Under the above circumstances, we might say that exchanging ‘Mary’ for ‘Mimi’, or vice versa, in the context (3) *induces* change in the proper understanding of that context. ‘Mimi’, for example, makes it much more likely that (3) bears the pet-keeping sense. In fact, that that name was used provides a very good reason for so understanding it. This phenomenon also points a moral for what is at stake in issues of s-opacity. Suppose that X \_\_\_\_ Y is s-transparent, and A and B are two names which, on your lips, would name the same thing/person. It does *not* follow that in a situation where you would (did) express a truth in the words XAY, you could just as easily have done so in the words XBY. For the truth of XAY does not guarantee that it will be easy to *say* XBY so as thereby to say the right thing. Under the circumstances, it might be quite difficult for you to produce XBY in such a way that the words X \_\_\_\_ Y would be to be understood as making that very contribution which they were understood to make when you said XAY. Circumstances plus choice of a name may have powerful effects on the proper understanding of other words with which the name is used—e.g., on what one *would* be talking about in just those words. Where this is so, it may contribute to an appearance of s-opacity where there is none.

Pseudo-opacity makes problems not just for recognizing, but also for deciding what is opaque and what is not. Consider the following variant on (2):

4. \_\_\_\_ is called *that* because of his size,

where ‘that’ refers to what goes in the blank. Again, substituting ‘Barbarelli’ for ‘Giorgione’ changes truth to falsity. But is this opacity? Note that the referent of ‘that’ has also been changed by the substitution. Does this respect the principle of constancy? Not if we count the referent as part of the contribution that words such as (4) make to what is said in them.

So what is to be done? First of all, that depends on what we ‘always wanted’ opacity to be, and I do not think *that* question has a clear and univocal answer. But suppose we decide that we want opacity to be something such that (4) certainly has it. Then there are two ways, at least, to achieve the effect. One is to explain further what we will understand by ‘contribution to what is said’.

In the sense that matters, one might hold, (4) *does* make the same contribution to what is said in both cases above: though the referent of ‘that’ and hence truth value vary with the substitution, the ‘that’ is to be understood in the same way in both cases with respect to the direction it gives on what the referent is to be, viz., whatever name or expression goes in the blank.

A better idea, however, is this. Note that with the pseudo-opaque (3), choice of the name ‘Mary’ rather than ‘Mimi’ induces but does not force a particular understanding of the words. We could construct circumstances where ‘Mary had a little lamb’ was spoken, referring to MacAdam and bearing the pet-keeping sense, even if this is unlikely to happen. No property of (3) itself, or a proper understanding of it—on any of its understandings—*requires* a change in what is said when we substitute ‘Mary’ for ‘Mimi’. In contrast, change in the contribution of (4) to what is said to be so is automatic and mandatory where one name is exchanged for another. Either the proper understanding of ‘that’ changes, or the referent does, with corresponding change in what is required for truth. This contrast suggests what I will call the *lenient principle of constancy* (since it allows more to be opaque): For any pair, (XAY, XBY), to demonstrate opacity, the contribution of X \_\_\_\_ Y to what is said must be the same in both members except where some property of X \_\_\_\_ Y, on a relevant understanding, forces variations from one member to another. On the lenient principle, (4) appears opaque. But it is also significant that there are such degrees of freedom in fixing the notion we want opacity to be.

#### IV

*Speaking variability.* The interest of pseudo-opacity depends largely on the kinds of possibilities there are for it to arise, and how easily recognizable these would be where realized. Here, a picture of what is said may stand in the way of seeing the interest. Philosophers often talk as if what is said in words, X \_\_\_\_ Y, depends only on what they mean (in English, in the case of (3) and (4)), and a handful of other easily specifiable factors, such as the time and place of speaking, and the referents of any clearly referential parts. If that is so, then as a rule pseudo-opacity should be easy to check for. In my view, however, the norm is that this is not so—particularly not in philosophically interesting cases. Here there is no space to argue the thesis in detail.<sup>2</sup> But by illustrating it I hope to show that, if it is right, then, given our present state of knowledge, issues of opacity are typically much cloudier than they are taken to be.

To begin with the simplest case, I borrow and expand on an example of Austin’s. An ichthyological expedition is working around the coral reefs of the

<sup>2</sup> For more than an indication of what is going on here, see especially my 1981*c* and Ch. 1 above.

Red Sea. Swimming around the reefs are many extravagantly coloured fish: bright red, blue-green, yellow and purple, etc. Specimens are taken. But when these are laid out on the deck, they all look muddy grey. On the other side of the boat, the cook is fishing for dinner. His prey are difficult to spot and spear, since, muddy grey in the open water, they seem like shadows. He does spear a large one, however, which, laid out on the deck, turns out to look distinctly dark red. Now consider:

5. There's a red fish on the deck. Bring it to me.

Such might be said to his assistant by a scientist wanting to make a last observation. It also might be said to his assistant by the cook, wanting to make dinner. Consider the first case. Suppose that, by chance, no specimen of the bright red species was taken. Then, I think, what the scientist said was false. The assistant cannot comply with the order by bringing the cook's catch. Now for the second case. Suppose the cook's catch had slid overboard. Then what he said was false, and the assistant could not comply by bringing one of the specimens, regardless of its species.

The point, without argument, is this. In each case, the word 'red' was used to mean *red*, and to speak of a fish's being coloured red. What varies from case to case is what would count as (a fish's) being coloured that colour. No range of ambiguities of 'red' in English could account for such differences in what would so count. So, the words 'is red', on a given meaning, may make various distinct contributions to what is said in speaking them, depending on what is to be taken to count as being as thus said to be, which is a matter of how the words, as spoken, were to be understood. An adequate specification of such an understanding is, then, something over and above a specification of what the words mean in English. Crucially, sameness of contribution to what is said, for two distinct speakings, is not guaranteed by constancy of meaning. In the case of 'red', it requires at least constancy of understanding of what being coloured red is to be understood to come to. But the point about the relation of meaning to contribution to what is said is, I think, absolutely general. What may need specifying to fix a contribution depends on the form of words involved.

With the general point in mind, consider one philosophically interesting example. Suppose Octave to have tried and failed to lift a weight. Now consider:

6. Octave could have lifted that weight.

When would (6) be true? A definite and direct answer to that question, if available, might be thought a contribution to philosophical discussions of human ability, or freedom, or responsibility. But now consider possible continuations of (6): 'if he hadn't ploughed the back 40 first that morning'; 'if he had been a better eater while growing'; 'if he had begun the Weider muscle development programme two months ago'. Reflection on cases where one continuation is

appropriate while another is not suggests, at least, that there are many distinct things to be said of Octave and that weight in some speaking of the words (6). Some of these may be true, while others are false. Where one of them in particular is said, a proper understanding, and a proper specification of it, depends on more than what we know about (6) so far, though we have not yet seen how to say what more is required. If this is right, then a philosopher who demands to know whether Octave ‘really’ could have done it, or whether (6), as thus specified, is true given such-and-such states of affairs, is simply making a mistake—one likely to engender philosophical problems of the sort which commonly arise when variety of the above sort goes unrecognized.

To know whether belief ascriptions are opaque, or merely give an appearance of being so, we must first know by what features one such may differ from another. Only then can we be in a position to check with confidence for pseudo-opacity. If there are many distinct ascriptions to be made to a given person in given words with meanings fixed, it will be important to know what, besides the meanings of the words, distinguishes them. About such matters I think we have little reason for confidence at present.

## V

*Speaker responsibility.* I will now present some reasons why belief ascriptions *should* be transparent, but *not* a demonstration that they are. Let us begin with e-opacity, and consider things of form (7), where A is a name:

7. S believes that XAY.

Recall that, if (7) is existentially opaque, this does *not* mean that something in that form of words might be true where the name A referred to nothing. But could such a thing happen anyway?

Consider (8):

8. Marvin went to the races.

Could (8) be true if ‘Marvin’ refers to no one? At first approximation, no. To see whether (7) and (8) differ in this respect, we must know why this is so. Here is a thought: (8) cannot be true where ‘Marvin’ refers to nothing, since in that case, it does not express a thought. For if it did, the thought it would express, or what it would say, was that such-and-such was so. What it would say to be so would be that someone or other went to the races. But then it is proper to ask, ‘*Who* was said to have gone to the races?’ Such a question can have no answer save, ‘No one’. And if no one was said to have gone to the races, it is difficult to see what *else* might have been said to be so.

If this thought is accepted, then there is a quick argument to a conclusion on (7). Consider, for example, (9):



## 9. Serge believes that Marvin went to the races.

If a thought is expressed in (9) at all, then it must be that Serge believes something or other. And it must be in the words (8), as occurring in (9), that it is said what Serge believes. But by hypothesis, (8) does not express a thought if 'Marvin' refers to no one. But if (8) does not express a thought, it does not express anything that anyone could believe, nor that could be true or false. Hence, in (8) in (9) it has not been said what Serge is being said to believe. Hence, where 'Marvin' refers to no one, (9) does not express a thought, hence cannot be true. (If (8) does not express something that *could* be believed, then (9) is certainly not *true*, in any case.) So no instance of (7) is true where the name A in it has no referent.

For some this will seem reason for rejecting the above thought about thoughts. Note, though, that the proper use of 'thought' is not what is at issue here. The operative principle here is: what you believe, as reportable in words of form (7), is something to be so. If someone *says* something to be so, you can believe it or not. But if there is nothing one has said to be so, there is nothing for you to believe or not. If this does not convince you that (7) cannot be true where A refers to nothing, then let us return to (8). There, we have already noted the general propriety of 'who' or 'which Marvin' questions. What this propriety points to, I think, is this: where a speaker purports to be telling us something, and he uses a name in doing so, as in (8), there is, *ceteris paribus*, a presumption that he knows who he is talking about. If he does, then we expect that he can *say* who it is, at least in vague terms like 'A guy I know', or at least that he knows that such specification could be given. If he knows that such specification could *not* be given, then he ought not to speak in those words. The bizarreness of doing so is recognizable. There is, apparently, some general stricture on the use of names, something like 'Be informed: Know who you are talking about in using them.' Suppose a speaker knows that 'Marvin', on his lips, would refer to no one. Then he knows that no specification could be given of who he was speaking of. So he ought not to say (8). Suppose you ask: 'But suppose he does say it anyway. Mightn't it be true?' Then a reasonable reply would be, 'Mightn't *what* be true?'

When we move from (8) to (9), we may observe that the stricture seems *equally* applicable in both cases. For the question 'Who did you say Serge believes went to the races?' feels just as legitimate and proper as the question 'Who did you say went to the races?' The thought is: suppose, in saying (9), you say something to be so. The only plausible candidate for what it is, is that Serge believes something to be so. But then it must be in the words (8), occurring in (9), that you say what this something is. But if there is nothing that you would say to be so here in saying (8), there can be nothing that you have said the something to be in saying (8) in saying (9).

If you know that 'Marvin' on your lips refers to no one, then, you ought not to say (9) equally as you ought not to say (8). You would be violating the stricture in doing so. And the question with respect to (9), 'Nevertheless if one

did say it, mightn't it be true?', equally invites the response, 'Mightn't *what* be true?' Our practice, at least, suggests that we recognize that we ought not to say such things. Suppose, e.g., Serge delusively thinks there is an extra member of the department named 'Marvin'. Then *he* might say (8), taking himself to say what he believes. But—at least without a good deal of story first—we would not say what we take Serge to believe by speaking (9). *Inter alia*, to do so would be to invite the wrong questions. The norm would be to seek some circumlocution, such as 'Serge believes there's this member of the department named "Marvin" and he went to the races'.

The force of the case here, and the intuition behind it, is easily overlooked if one is already convinced that there obviously are true instances of (7) with a name, A, referring to nothing—as many philosophers appear to be. But what is the evidence for this? I think philosophers tend to think of examples like

10. Mimi believes that Santa Claus lives at the North Pole

11. Octave believes that Pegasus had supernatural powers

and speakings of these in which truths obviously are expressed. The trouble with this evidence is that in the special sorts of discourse in which (10) and (11) are alright, *any* discursive form of words might express a truth while containing a name referring to nothing—as evidenced by ways (10) and (11) might be continued. Suppose Mimi is a little girl and Octave a classical scholar. Then, in appropriate situations, (10) and (11) might be followed respectively by (12) and (13):

12. And she's right—that *is* where Santa Claus lives; not in Hawaii as little Didier thinks.

13. But he didn't. Pegasus was a perfectly ordinary *winged* flying horse.

Where (9) (10) might express a truth, then, so might (12).<sup>3</sup> 'Who' or 'which' questions are equally pertinent for (10) as for (12), no more and no less. If existential generalization fails in *this* way for belief ascriptions, then it fails for *everything*—reason to think that this sort of evidence sends us sniffing down the wrong track.

The foregoing does not settle whether belief ascriptions are e-transparent. If Quine's argument is good, that all depends on whether words of form (14)

14. S believes that X \_\_\_\_ Y,

the blank fillable by a name, may be read as ascribing a property to something. On the argument, that requires their being s-transparent. So what are the chances

<sup>3</sup> One insider to another: 'Well, what did Marvin do today?' If you like solemnity, imagine Serge's tales achieving literary importance.

that they are that? Suppose one MacNab to go by two names, 'Grant' and 'Gustaf'. Now consider:

15. Brangwyn believes that Grant was at Harrod's.
16. Brangwyn believes that Gustaf was at Harrod's.

Again, to some philosophers, it has seemed obvious that, e.g., (15) might be true while (16) is false. What they think of are cases where Brangwyn knows MacNab to be named 'Grant' but does not know him to be named 'Gustaf'. But to clarify the issue: Of course, some *speaking* of (15) might be true while some speaking of (16) was false. But could a speaking of (15) be true while a speaking of (16) *related to it by respect for the (lenient) principle of constancy* was false?

There are signs, at least, that the answer to this last question is 'no'. First, to adumbrate a later issue, consider (17):

17. If you ask Brangwyn, 'Was Grant at Harrod's?', she will say yes, but if you ask her, 'Was Gustaf at Harrod's?', she will say no.

Unproblematically, I think, it might be correct to say this, and it might even be true. Now consider (18):

18. Brangwyn believes that Grant was at Harrod's but she does not believe that Gustaf was at Harrod's.

If belief ascriptions are truly s-opaque, one might expect that (18) would be as unproblematic as (17). If we reflect carefully, however, I think we will see that if we (audience included) know that Gustaf is Grant is MacNab, then (18) is not something we would ever say (outside of philosophy perhaps). What is wrong with (18)? One thing might simply be this: it is abnormal to use two different names in such close proximity to name the *same* person. If this is so, it is not such a superficial fact as it may seem. After all, if it is just a plain truth that Brangwyn believes that Grant was at Harrod's, and equally a plain truth that she does not believe that Gustaf was at Harrod's, why should we not be able to say so simply, in the most obvious way? But further, if you are spoken to in words (15), a natural way of summarizing the information is: 'So she thinks he was *there*, does she?' Now suppose the first barrier is passed and we do take 'Grant' and 'Gustaf' in (18) as referring to the same person. Then, on hearing the second half of (18), it is difficult to avoid the feeling, 'But I thought you just said she *did* think he was there'.

Both these signs, I think, have a deeper root. Again, it has nothing specifically to do with belief ascriptions. Consider (19):

19. MacNab was at Harrod's.

Suppose we are speaking to someone who does not know that MacNab is named 'MacNab'. Then normally, even if we take it to express a truth, we will not

speak to him in words (19), at least not without first telling some story so as to introduce the name properly. We might think of this as respect for an audience's (expected) knowledge: choose names that your audience will take aright. But then, suppose there is a fact statable in (15), but simply not statable in (16) at all. Now suppose we are speaking to someone who knows MacNab as 'Grant', but not as 'Gustaf'. Then this is a fact we cannot tell him without violating the above stricture (or engaging in circumlocution). If belief ascriptions really were *s-opaque*, it would sometimes be difficult to say what was so (and impossible to communicate it) while observing normal rules of conversation. In fact, we do not *seem* to be so encumbered. *If* it is accepted that what is said in belief ascriptions does not thus conflict with normal strictures on conversation, then this has profound consequences for analyses of such ascriptions, as we shall soon see.

## VI

*On what belief ascriptions say.* Whether belief ascriptions are opaque, if there is a fact of the matter, is not easily read off the phenomena. What is not obvious is what needs preserving to respect a principle of constancy. For that all depends on how what is said in one belief ascription may differ from what is said in another. Which depends, in turn, on what a proper understanding of such an ascription consists in. I think there needn't be any unique right way of saying what is said in such ascriptions—some 'logical form' which just *is* theirs, as it were. If not, then that is *one* cause for worry as to whether there is a fact of the matter here. What follows, however, is a suggestion on one way of saying what belief ascriptions say.

We already noted a case of which there was something true and something false to be said in given belief-ascribing words, (16). Apparently, then, there is a variety of things to be said of a given situation in one such form of words. What differences make for the variety? Let us introduce the term *consense* for 'conventional sense and reference'—i.e., whatever is fixed about what is said by what the words used mean in their language (on relevant senses) and the identities of any items or individuals referred to. Now it will be especially difficult to see where constancy is maintained if there is a variety of things to be said in given belief-ascribing words with given consense—i.e., if belief ascriptions exhibit what I have termed speaking variability. Let us now investigate that possibility.

If there is speaking variability for words of form (7), there are two relevant places where the differences might come: in the words, 'that XAY', or in 'believes'. Consider the first possibility first. The most natural line for developing it would be to take 'that XAY' at face value as a referring expression. What it would refer to is that which is believed. Call such an item, whatever it may be, a thought. Then different speakings of the same words, 'that XAY', may refer to different thoughts. Given such variation, one speaking of (16) may be true while another

is false, because in each a different thought is referred to—one that Brangwyn believes in one case, and one that she does not believe in the other.

If this is the right story and the whole one, then, with consense fixed, variation in belief ascriptions is variation in what a believer is said to believe. Fix something to be believed, and there is one state of a believer which could make it true that he believed that. Ascriptions of belief in one such item, then, will be true or false *tout court*. The story is congenial if we think of the function of belief ascriptions as recreating a bit of how the world looks from the believer's point of view.<sup>4</sup> And no doubt there *are* many subtle differences between things one can be said to believe. But just there lies a problem. For if telling the truth about someone depends on picking out from this assortment just the right item which is that which he believes, then telling the truth about belief may be none too easy a thing to do. To do so, one will have to get one's words to refer to just the right thought—'the one the believer has', if there is such a thing. But it may be difficult to find the circumstances for doing that.

Consider this case: Rudolf, studying in Cambridge, gets a letter from his mother in Linz. In it is (20):

20. Dr Lauben was wounded.

Rudolf thinks of his childhood physician. Later he visits the hospital in Cambridge. There he sees a heavily bandaged man in a bed with a name card, 'Dr Lauben'. He takes the man to be called 'Dr Lauben'. He does not take him to be the *same* Dr Lauben as the one his mother wrote about. But it is one and the same man.

Rudolf *might* disbelieve his mother but believe what he sees in the hospital, or vice versa. In such a case does he believe that Dr Lauben was wounded? There appear to be both true and false things to be said in saying so, though there will be different truths to tell, in each of the above two events. On the present analysis, such differences must turn on differences in what Rudolf believes: where he disbelieves the letter but thinks the hospital scene no sham, there is something which he believes, expressible in (20), and you speak truth in saying him to believe that Dr Lauben was wounded when, but only when, your words 'that Dr Lauben was wounded' refer to that thing. In the reverse situation, there is also something Rudolf believes, expressible in (20), for which the same remarks apply. But in this case, it is a different thought that he believes; you must make different references tell the truth.

How do you go about making the right references here? Suppose that we, in Linz, know that Dr Lauben is playing an elaborate hoax. He has bribed Rudolf's mother to write the letter. Then, just to make sure, he went to Cambridge, where he had himself bandaged and put in a hospital bed. And he has yet further tricks

<sup>4</sup> As far as I can see, many philosophers do take such a view of belief ascriptions. An example would be Dummett 1981, esp. pp. 108–13.

up his sleeve. What he wants is Rudolf's sympathy; if he gets it, he wins a bet. In the hospital, Rudolf reacted in the right way. Now someone comes to us who has the news. Suppose he wants to tell it by saying 'Rudolf believes that Dr Lauben was wounded'. Then, to tell us the truth, he must get us to understand the words 'that Dr Lauben was wounded' as picking out just 'the thought that Rudolf had' under the above described conditions. Perhaps with some stage setting he could do this. And we *might* be grateful for the extra information we would thus get. But it does not seem to be mandatory for him to do this in order to tell us the truth in the above words—the truth that counts the most, in fact, for our purposes.

The foregoing is an indication, though no more, that the above analysis has got the function of belief ascriptions slightly wrong. Perhaps they are not meant (just) to re-create a bit of the believer's view of the world. The alternative view would be that they function to relate the believer in certain ways to the speaker of a belief ascription and his intended audience. Suppose there is some thought that *we* are interested in—the one that *we* would express in given circumstances in speaking (20). Then, on this view, saying 'Rudolf believes that Dr Lauben was wounded' may be a way of relating Rudolf to *that*. We may ask, 'Does Rudolf believe that Dr Lauben was wounded?' Depending on what we thus express, and the facts about Rudolf, the answer may be, 'in a way, yes, and in a way, no'—an indication that there are both true and false things to be said in saying him to believe *that*. Which is both a natural way of describing the variety in cases like the above and, on its face, prejudicial to the original analysis. If this is accepted, then we must reject the view that, for a given item to be believed or not, there is some *one* particular state of a believer which could make it true to say that he believes that, and hence the view that with such an item and consense fixed, 'S believes that XAY' will be either true or false *tout court*. In that case, we must look for variety in proper understandings of 'believes' on various speakings of it.

If we admit that 'believes' may make different contributions to what is said in different speakings of it, then we have already made an important break with the picture previously described. There may still be a variety of thoughts to be expressed in given words, e.g., (20), hence a variety of things one may be said to believe, even in words with fixed consense. What there will not be within this variety is some one thought which is 'that thought which the believer *really* believes'. For there *is* no reality corresponding to *that* notion. On the right occasion, in the right speaking, to be understood in the right way, a believer such as Rudolf or Brangwyn might correctly be said to believe any of many or perhaps *any* of the thoughts within such a variety, depending on precisely what relations he or she is thereby said to bear to it, or what is to be taken to count as his or her doing so. Again, the emphasis is shifted, on this account, from how the believer would express, formulate, or picture things to how *we* would express things, and how the believer relates to *that*. In a belief ascription, we are, as it were, fitting the

believer into *our* picture of the world, not fitting ourselves behind his 'window on it'.

Fix what is said to be believed in given words, 'that XAY'. Then such a variety in understandings of 'believes' may be thought of as variation, from speaking to speaking, in what believing that is to be taken to come to. To see how to describe it, we might ask: what is it to believe that XAY? When would someone do that? As an answer, philosophers sometimes seem to accept what might be called an assent test. Brangwyn believes that Grant was at Harrod's, for example, roughly just in case she is prepared to assent to the *words*, 'Grant was at Harrod's'. If that were exactly right, then there would be an easy proof of the s-opacity of belief ascriptions. For of course she might assent to words 'Grant was at Harrod's' but not to words 'Gustaf was at Harrod's'.

For a variety of reasons, however, such an answer cannot be quite right, at least for any given assent test. One set of reasons concerns what might be called blocks to assent. Leibniz, Freud, and Chomsky, for example, have each presented us with cases of different sorts where it would be natural to say that someone believes such-and-such, where it is not demonstrably wrong to say this, but where, for one reason or another, the believer would not assent to words in which the ascription was naturally put. A Leibnizian example: wherever someone recognizes a truth of the form 'If P then Q', and recognizes the truth of the relevant P, he is prepared to conclude Q. We say that he recognizes the validity of *modus ponens*. We might also say: he believes that *modus ponens* is valid. Depending on the circumstances of our speaking, we need not be daunted on discovering that the person has never heard of the term '*modus ponens*'.

Even if such considerations are not convincing, there is a second set of problems with assent tests. Briefly, no such test cuts fine enough. Consider Rudolf. *Sometimes* he would assent to words (20) and sometimes he would not, depending on how he understood those words. For that matter, the same is true of Brangwyn: she will assent to words, 'Grant was at Harrod's' perhaps, but only where she takes 'Grant' and 'Harrod's' as referring in the right way, and 'was' as picking out the right time. What does such selective assent show? If we accept that there is a variety of things to be said in saying Rudolf to believe (20), then the most it can show is the truth of some of these and/or the falsity of others. So at the least more needs to be said about what the truth of belief ascriptions depends on.

Before dismissing the idea of assent tests, however, let us note what they require. We can distinguish two components. First, there is something which a believer is required to do to *evince* his belief. Second, there are circumstances under which he is supposed to do it, namely, when it is *elicited* from him in a certain allowable way. Problems come when we suppose that there is just one thing that evincing might come to—assent to indicated formulations of the relevant thought, say—and only one form of elicitation that could be allowable—asking a properly formulated question, say. Instead of supposing

that, though, we might try supposing that contributions of 'believes' to what is said in speaking it differ from one another precisely in terms of what sort of evincing would be pertinent to showing that the believer believed what he was said to, and what sort of elicitation would be right for yielding this evincing. A given understanding of speaking of 'believes' might then be specified by specifying what evincing is required and what elicitation permissible if things are to be as said to be in speaking it. On some such understanding, for example, Brangwyn's saying yes to a question 'Was Grant at Harrod's?' *might* show things to be as said to be in some speaking of (16), if the evincing required on *that* understanding of 'believes' allows for this.

I will not defend this suggestion further here. Taking it as a sample of what an account of speaking variability might look like, I turn now to examining its consequences for issues of opacity. I begin with s-opacity. Suppose that there could be a true speaking, S, of (15), such that there could be a speaking, T, of (16), which maintains required constancy with respect to S, but which is false. That is to say, the change from truth to falsity must be achieved solely by exchanging the name 'Grant' for 'Gustaf', plus whatever other changes that change might require.<sup>5</sup> First, then, Brangwyn must fail in the required evincing towards what is said in T in 'Gustaf was at Harrod's'. This means that evincing which was adequate for the truth of S is not adequate for the truth of T. It is hard to see how this could come about unless the required evincing in the case of T somehow essentially involves the name 'Gustaf' — for example, perhaps she must say yes when asked, in those words, 'Was Gustaf at Harrod's?' But second, 'believes' must bear the same understanding on its use in T as it does on its use in S, except where forced to do otherwise. If such understandings are to be specified in terms of evincing, then there must be some sense in which the same evincing is required in the case of T as in that of S.

For both these requirements to be satisfied at once, it must be that the understood requirement on evincing in the case of S needs specifying in such a way that when it is applied to the facts in the course of evaluating S, it yields one set of results, but when applied to the same facts in the course of evaluating T, it yields another. For example, the requirement in the case of S might be that the evincing essentially involve whatever name comes first to the right of the 'that' in words for which it is understood to apply. Apply this requirement when considering S, and it will be satisfied if Brangwyn behaves appropriately *vis-à-vis* 'Grant'. Apply it while considering T, and she must behave appropriately *vis-à-vis* 'Gustaf'.

The above, of course, is not the only plan for saying that appropriate evincing must essentially involve a name. One might also say simply: In the case of S,

<sup>5</sup> It is perhaps worth emphasizing: the exchange of names is in what is said, or better, in what is spoken; not in specifications of proper understandings of what is said. The understanding of everything outside of the name place must remain constant, though if some such thing is understood as referring to the name in that place, then constant understanding plus different name *may* yield different results.



Brangwyn's evincing must involve the name 'Grant'. If *that* is the requirement, then it will be satisfied in the case of T just in case Brangwyn evinces belief in what is said in T in 'Gustaf was at Harrod's' when it is formulated using the name 'Grant'. This, of course, she will do, given the truth of S. On such an understanding of 'believes', if there is one, then, (15) is *s-transparent*.

On the present analysis, then, what is needed to show that some belief ascriptions are opaque is an argument that there are possible understandings of speakings of 'believes' on which the evincing required, given what believing is to be understood to come to, can only be specified in the first way and not in the second. It is hard to see what such an argument would be. Suppose we can find none. Then one possibility is that where, if ever, an understanding of 'believes' makes reference to some specific name, this understanding is always to be specified in the second and not the first of the above ways. (An indication: suppose you are saying something of form (7), intending to use one name—'Gustaf', say—but absentmindedly using another—'Grant', say. Are there ever circumstances where you would actually have to take back what you said? I doubt it.) If this possibility is realized, then all belief ascriptions, even where the understood evincing requirement essentially involves a name, are *s-transparent*.

A second possibility is that although, perhaps, we sometimes understand 'believes' in such a way that, for what believing is to come to, evincing must essentially involve some particular name, nothing about such understandings shows that they must be specified in one of the above ways rather than the other. For some speaking of (15), perhaps, we expect that, if what was said is true, then we can successfully interrogate Brangwyn using the name 'Grant'. But there is nothing to show that our expectation is one 'that we can use *that* name', 'that' referring to whatever occurred in the appropriate place in (15), *rather* than an expectation that we can use 'Grant'. If this is the way things turn out, then our notion of opacity, such as it is, begins to dissolve: there are things to be said in belief ascriptions, on some understandings of them, such that when those things are specified in one way the ascriptions *look* opaque, and when those same things are specified in another, the same ascriptions look transparent. Beyond that, there will be no fact of the matter as to the opacity of such belief ascriptions. My guess is that if there are any such understandings, this is the closest any belief ascriptions get to being *s-opaque*.

Let us now turn to *e-opacity*. If a speaking, S, of (15) is true, does that guarantee the truth of (21)?

21. There is someone such that Brangwyn believes that he was at Harrod's.

First, (21) could only be true if there is something which it says to be so. We are supposing that that could only be so if there is some property which (22), on a speaking, ascribes.

22. Brangwyn believes that \_\_\_\_ was at Harrod's.

If not, it is not just that (21) *might* not be true where (15) is; (21) *cannot* be true whether (15) is true or not. Second, if there is a variety of things to be said in (15), there may be a variety of things to be said in (21) as well. If so, what needs guaranteeing is the truth of *some* such thing—one related appropriately to what was said in S. Such a thing would be said in a speaking of (21) in which all parts of it occurring in (15) are to be understood in the same way as they were to be on S, insofar as this is possible.

Turning to the first issue, does (22), on a speaking, attribute a property? To all appearances, (21), on a speaking, says what might be true. So to all appearances, yes. The next question is: What (sort of) property? Note that if S is true, ‘Grant’ referring to MacNab, there is surely *some* property which MacNab is thereby guaranteed to have: namely, the property of having a thought about him, to the effect that he was at Harrod’s such that Brangwyn stands in a given relation to that thought: namely, the evincing–elicitation relation indicated by a proper understanding of ‘believes’ in S. Could this be the property ascribed on an appropriately related speaking of (22)? If so, then the truth of S *would* seem to guarantee the truth of what is said in appropriate speakings of (21). So belief ascriptions would be e-transparent.

Here, I think, conventional wisdom plays a trick on us. On that wisdom, there are two senses of the English words (15), an s-transparent one and an s-opaque one. The s-transparent one may be read roughly as (23):

23. Grant is such that Brangwyn believes of him that he was at Harrod’s.

The s-opaque sense *must* be e-opaque, by Quine’s argument. The s-transparent sense *may* be e-transparent, and if (23) is the right reading of it, apparently will be. But now, a problem: with consense plus an understanding of ‘believes’ fixed, the truth of (15) on its supposed s-opaque sense certainly guarantees the truth of (23). If it were not enough to do so, and if conventional wisdom were correct (e.g., about assent tests), then (23) could *never* be true: surely there is *some* way of referring to MacNab such that Brangwyn would not recognize him so referred to. But the truth of (23) certainly guarantees the truth of (21). But then so does (presumably) s-opaque (15). So why isn’t (15) just e-transparent on any of its senses? The answer, if there is one, must be that the version of (21) whose truth is guaranteed by that of (23) is not appropriately related to the supposed s-opaque version of (15). What *that* would have to guarantee in order to be e-transparent would be the truth of (21) in some other sense of those words—a sense which seems not to exist, not even as a piece of nonsense like ‘someone was so-called because of his size’. Such *might* have been taken as an argument against splitting senses in the conventional way. Instead, it seems to have been taken as reason for letting the supposition stand that there is an s-opaque sense of belief-ascribing words on which they are also e-opaque.

The present viewpoint, in contrast with the above, is that, whether or not there are some belief ascriptions which are s-opaque, the phenomena can be

accounted for with one form of analysis, and without positing any *ambiguity* in belief-ascribing words. Words (15), for example, have one *form* of truth condition on the present view. If some speakings of (15) say what is true while others say what is false, all speaking of one given state of affairs, that is because there are various things that *believing* something may sometimes correctly be taken to come to, and correspondingly different things that may be said in *saying* someone to believe something, on various understandings of 'believes'. What is at stake is speaking variability, not ambiguity in English. If that is correct for (15), then it is also to be expected that there can be one form of truth condition for (21), even if there are various things to be said in it on given consense. Here is a first stab at an analysis: in (21) on a speaking, 'he was at Harrod's' indicates a thought; 'believes' indicates some particular evincing–elicitation relation, R, and what was said in (21) is true just in case there is someone such that Brangwyn stands in R to that thought about him. If that analysis is right, then, starting with a speaking, S, of (15), and maintaining the understandings of all its parts occurring in (21), we get an understanding of (21) on which its truth would appear to be guaranteed by the truth of S. So belief ascriptions appear e-transparent.

To this attempt, one might object: 'he was at Harrod's', in (21), does not express a thought, since there is no one to whom 'he' refers. But now consider (24):

24. There is someone such that he has a beard.

'He has a beard', in (24), equally fails to express a thought. So how could (24) say what is true? The answer is a general remark on how existential generalizations are to be evaluated. We will count (24) as true if we can find someone such that if we suppose the 'he' in (24) to refer to him, then, on that supposition, we get a true thought. Suppose we treat the suggestion about (21) in the same spirit. (21), on the analysis, should be counted as true if we can find someone such that, on the supposition that the 'he' in (21) refers to him, the thought expressed in 'Brangwyn believes that he was at Harrod's' is true. This will be so just in case Brangwyn stands in the right evincing–elicitation relation to the thought expressed *on the supposition* in the words 'he was at Harrod's' in (21).

But on this scheme of things, couldn't there be some belief ascriptions which were e-opaque? Such would happen for a speaking, S, of (15), if Brangwyn stood in the indicated relation, R, to the thought expressed in S in 'Grant was at Harrod's', but did not stand in that relation to the thought expressed on the supposition in 'he was at Harrod's' (choosing MacNab as the man for 'he' to refer to). We are already familiar with the sort of understanding of 'believes' that would be involved here, if such exists: an understanding, e.g., that was only correctly specified by describing the required evincing as essentially involving whatever name appears in a certain place, e.g., to the right of the 'that' in the words in question. For the same understanding of 'believes' as occurring in (21), such a specification would pick out nothing: there is no such name. But if the name essentially involved in evincing, for a given understanding of 'believes',

could be specified in the above way, what would be lost by specifying it as follows: ‘whatever name occurs to the right of “that”, if there is one, and if not, whatever name would occur if the truth were to be told exactly’? On such a specification of an understanding, (21) is guaranteed to tell a truth if (15) does. And there would seem to be no fact about a speaking of (15) which would show this last specification incorrect where the trouble-making specification, above, *was* correct. I conclude that, insofar as there are facts of the matter about such things, belief ascriptions are e-transparent.

## VII

*Quine’s argument revisited.* In the (unlikely) event that belief ascriptions turn out to be e-transparent and s-opaque, rather than transparent, full stop, we had better look at Quine’s argument again. Here is a schematic suggestion on how the argument could break down. Suppose there is an open sentence, ‘X \_\_\_\_ Y’, which works like this. First, there is some definite property—call it P—which it ‘indicates’ or ‘speaks of’. Second, however, filling in the open space with a name may change the understanding of ‘X \_\_\_\_ Y’, so that it speaks of a different property, the property depending on the name used. If the name is A, call this property  $P_A$ . Third, any property  $P_A$  is a specialization of P in the sense that anything which is  $P_A$  is (and must be) P as well. Now suppose A and B to be names, both of which refer to an object O. Then the situation is this: ‘XAY’ will be true iff O has  $P_A$ ; ‘XBY’ is true iff O has  $P_B$ ; and ‘There is something such that X if Y’ will be true iff something has P. Conceivably, O may have  $P_A$  but lack  $P_B$ . If this is possible, then ‘X \_\_\_\_ Y’ is referentially opaque. But if O has either  $P_A$  or  $P_B$ , then O has P. Hence something has P, hence the last existential generalization is true. So ‘X \_\_\_\_ Y’ is existentially transparent. For such an open sentence, Quine’s argument fails.

Might one think of belief ascriptions on the above model? Intuitively, this idea is no less plausible than the idea that some belief ascriptions are s-opaque. The ‘minimal’ property, P, attributed to O in ‘S believes that X \_\_\_\_ Y’—the one which makes for existential transparency—would be that of being such that there is a thought about it, one such that S stands in some belief-constituting relation to that thought. The specialized properties  $P_A$  which make for referential opacity would be ones which O has where there is such a thought, and the relation which S stands in to it essentially involves the name A.

## VIII

*Opacity and ‘direct reference’.* The supposed opacity of some belief ascriptions has been viewed as a test case for a certain view of names, attributed to Mill and

recently advanced by Kripke. We are now in a position to see why, even if such opacity exists, it is no test case at all. On the Mill–Kripke view, exchanging one name for another while preserving the referent, at least in simple cases, cannot change the thought expressed. So, taking all references as to MacNab, ‘Grant was at Harrod’s’ and ‘Gustaf was at Harrod’s’ must express the same thought. The idea of the test case, then, was this: suppose (15) might be true while (16) is false. This means that there must be something which Brangwyn believes which makes (15) true, and something which she does not which makes (16) false. But such could only be if the thought expressed in (15) in ‘Grant was at Harrod’s’ is different from the thought expressed in (16) in ‘Gustaf was at Harrod’s’. Thus are Mill and Kripke refuted.

I will make two remarks about this idea. First, there cannot be a relation, R, and a thought, T, such that Brangwyn stands in R to T and Brangwyn does not stand in R to T. So if ‘believes’ in belief ascriptions indicates a relation *and* if it always made the same contribution to what was said, then the above idea would be correct. If, however, ‘believes’ exhibits speaking variability, so that it makes different contributions to what is said in different speakings of it, then the case is not clear. I have indicated above that ‘believes’, on different speakings, may indicate different relations for Brangwyn to bear to something. Or, insofar as we view it as always indicating one relation (which is possible), there are different states of affairs, for different speakings, which would count as Brangwyn standing in that relation to something. If this speaking variability of ‘believes’ is adequate to account for the differences between true speakings of (15) and false speakings of (16), then the above idea of a test case is wrong. Differences in true value between some true speaking of (15) and some false speaking of (16) will be due to different contributions of ‘believes’ to what is said in each case, hence need not result from a different thought being spoken of each time.

The truth of some (15) and the falsity of (16), even given the Mill–Kripke view of names, need not commit us to the absurd view that, for some R and some T, Brangwyn bears R to T and she does not. Insofar as we may speak of *one* relation responsible for the truth of (15) and the falsity of (16), it is one that Brangwyn may sometimes be counted as bearing to something, while she is sometimes (correctly) counted as not. There is no obligation to count things in both ways at once.

The second remark is this: insofar as there are any convincing examples of opacity, they seem to depend heavily on saying what refers to the words used to say it. ‘\_\_\_\_\_ is so-called because of his size’ is about the best example we have seen so far. If there are understandings of (15) and (16) on which they are opaque, these would seem to adhere to the rule. The problem about such a speaking of (15) would be that Brangwyn must evince belief using the very name used in (15) to say what it is she believes. Depending on how this requirement is put, she might, perhaps, satisfy it for (15) while failing to satisfy ‘the same requirement’ for (16). But if that is what makes for opacity, then we hardly have an argument

against the Mill–Kripke view yet. The truth of (15) simply shows that there is some thought such that Brangwyn counts as believing it *when it is formulated in a given way*. Similarly, the falsity of (16) shows that there is some thought which she does not count as believing when it is formulated in a certain way. Such by itself shows nothing about the identity or distinctness of the thought in the first instance and that in the second.<sup>6</sup>

<sup>6</sup> I have been greatly helped in my thinking on this subject by many conversations with John Campbell. I would also like to thank the members of the ‘Friday Seminar’ in the rue Serpente, Paris, particularly Gilles Fauconnier, François Recanati, and Pierre Jacob.

# 9

## Vagueness, Observation, and Sorites

Does vagueness lead to inconsistency in language? Frege seems to have thought so, and Michael Dummett has suggested that he may have been right, at least in some cases. Dummett writes, ‘the use of vague predicates—at least when the source of the vagueness is the non-transitivity of a relation of non-discriminable difference—is intrinsically incoherent’ (Dummett 1978/1975). The moral Dummett draws is that ‘while our language certainly contains observational predicates as well as relational expressions, the former (though not the latter) infect it with inconsistency’ (p. 268). I will argue in what follows that Dummett is mistaken on both scores. First, our language does not and could not contain observational expressions in Dummett’s sense, nor is there reason to want it to. Second, our language is not and could not be inconsistent for anything like the reasons Dummett offers.

### 1. Content and Circumstances

Dummett appears to have a certain conception of what vagueness would be like. On it, first, it is (e.g., English) *predicates* which are either vague or not, and second, for a predicate to be vague is for there to be (possibly) some range of items such that nothing about what the predicate means determines whether it is true of the items or not. There are, on the conception, items of which the predicate clearly is true, full stop, and other items of which it clearly is not. But then there are yet other items for which it is not so, or at least not clear that the predicate is true of them, but equally not so, or unclear that it is not. Within this conception, trouble arises in particular where a collection of items may be arranged in a series so that the series begins with an item of which the predicate is clearly true, ends with one of which it clearly is not, and so that it becomes progressively less clear what to say as one moves towards the middle from either direction. So we may think of the area of application of the term as having fuzzy edges.

Before examining how vagueness might make for inconsistency, I want to make one point about this idea of what vagueness is. The point is that, while for any predicate in a language we could speak there may be items such that no facts about what the predicate means in the language determine whether it is true

of them, it is also so that what the meaning of a predicate does not determine may be determined by its contribution to what is said on a speaking of it. For predicates in general exhibit what I would term occasion-sensitivity: a proper understanding of a predicate on one speaking of it may differ from that which it bears on another speaking, where the differences are manifest, *inter alia*, in what the predicate as thus spoken would be true of.

Consider, for example, 'weighs 80 kg'. Suppose that 80 kg is Hugo's recommended weight. One morning, after months of dieting, he steps on the scale and *voilà*—80 kg. Later in the day, heavily dressed but without over-eating, he is such that should he step on a scale, it would register 84 kg. Is the predicate 'weighs 80 kg' true of Hugo at this time? I would argue that nothing about the meaning of that predicate determines this. The reasons, in outline, are as follows. First, weighing oneself in the morning is a generally accepted and recognized way of determining what one weighs, where weighing such-and-such is regarded as a semi-durable property, i.e., a property which one maintains over time barring over- or under-eating, etc. Normally, you do not count as gaining weight just by putting on clothes. Second, what you weigh is also generally regarded as a matter of the force you would exert on floors, scales, etc. And the force Hugo would exert at the time in question is 84 kg. Third, nothing about our concept of what it is to weigh 80 kg shows the first consideration to take precedence over the second, or vice versa. Note the general form of the situation here. There is a reason for saying that the predicate is true of Hugo and a reason for saying that it is not. Nothing about the meaning of the predicate shows one reason to outweigh the other. (Elsewhere (Ch. 1 above) I have argued that nothing about the meaning could do this if 'weighs 80 kg' is to speak of weighing 80 kg.)

Suppose, however, that someone says that Hugo weighs 80 kg, and says it for some purpose. For example, Hugo may say, 'I weigh 80 kg', announcing the success of his diet. One would not show what he said to be so not to be so by pointing out, 'Oh—with that overcoat and those boots?' On the other hand, suppose that Hugo wants to cross the trestle bridge across the ravine. It might be pointed out to him that the bridge is quite delicate and cannot bear more than 81 kg. Suppose he then replies, 'Oh well, I weigh 80 kg.' Given what he thus says to be so, he *might* be refuted in the above way. So there are various things to be said in saying, 'I weigh 80 kg', some perhaps true while others are false. If this is right, then, fourth: while facts about the meaning of 'weighs 80 kg' do not determine that this predicate is true of Hugo, or that it is not, facts about a speaking of that predicate may determine that what is said in it on that speaking is true of Hugo, or that it is not—depending on what those facts are. Such is the way of predicates.

This point about the sorts of facts that may determine applications is, I think, true of predicates in general. It will be useful in what follows. With the point as background, then, let us turn to examining the troubles that vagueness is supposed to cause.



## 2. Sorites Puzzles

The Sorites of legend concerns a heap. One grain of sand cannot make for heaphood or its lack. But then, the argument goes, neither can the whole heap of grains. How *would* the argument go? Look at a modern reclothing due to Hao Wang. Consider being a small number. 1 is surely small. In fact, since it cannot make much difference to size, if  $n$  is small, then so is  $n + 1$ , for any  $n$ . But then mathematical induction applies. Conclusion: all numbers are small. What makes *this* argument work are two apparently sound premisses (as for the second, just try to find a counterexample!) and induction. Now imagine numbering the grains in the heap, and consider the property of not taking away heaphood. Premisses of the same form, with the same principle, will yield like results.

For such puzzles, vagueness might be thought to do the following work. First, if a predicate  $P$  is vague, this might explain why, as we go through a series of items, it becomes less and less clear whether  $P$  applies without our being able to identify some definite point at which it ceases to apply. Second, one might hope to explain why, nevertheless, unclarity gradually shifts back into clarity, until, at last, there is at least something to which  $P$  does not apply. Third, one might hope, unclarity may explain the absence of clear counterexamples to the principle *if  $P(n)$  then  $P(n + 1)$* : it is never *much* less clear that  $P$  applies to  $n + 1$  than that it applies to  $n$ ; it is just that, as we move through the series, it gets pretty unclear of *both* adjacent members whether  $P$  applies. Fourth, one might hope to explain why we must either reject an eminently plausible principle of the form *if  $P(n)$  then  $P(n + 1)$* , or reject induction.

If we *must* reject either induction or a principle of the form *if  $P(n)$  then  $P(n + 1)$* , then it had better be the latter. Impressed by the lack of clear counterexample, e.g., in the above cases, one might think of looking at adjacent items in the middle of the series, and have the idea 'If  $n$  is  $P$  then so is  $n + 1$ . But  $n$  is not clearly  $P$ , thus neither is  $n + 1$ '. Saved by vagueness here. But what about the first item in the series? The principle takes us to the second, which we now have decided to call  $P$ . But that decided, the principle takes us to the third item, and so on. There seems no stopping. Which brings up a second point: for finite series, induction, while handy perhaps, is dispensable. Exhaustion will do. And there are versions enough of the Sorites, the heap being one, where the relevant series are finite.

If the idea of vagueness is to resolve Sorites paradoxes, it must explain why, where the paradoxes appear to arise, the relevant principle *if  $P(n)$  then  $P(n + 1)$*  is invalid. On its above application, it does not tend to do that. Worse, Dummett<sup>1</sup> and Crispin Wright (esp. 1975) cite instances of finite Sorites-type series where,

<sup>1</sup> Esp. Dummett, 1978/1975. A warning is in order: throughout when I speak about what Dummett and Wright have shown, I am not necessarily reporting what they take themselves to have

so they argue, the relevant principle, *if  $P(n)$  then  $P(n + 1)$* , cannot be rejected. If the vagueness of a predicate suggests otherwise, in such a case, then the suggestion *must* lead to inconsistency. A forceful example: it is possible to arrange a series of colour samples so that each two adjacent samples are indistinguishable in colour by looking (by human beings, in excellent light). Yet the first sample in the series may be easily distinguishable from the last—in fact, clearly red while the last is clearly orange. Physics knows finer distinctions than the human eye, so to speak—with the upshot, as they would put it, that indistinguishability of colour is non-transitive.

It is a natural thought that if two things are indistinguishable to the eye, *qua* colour, then they are the same colour. If A is colour C, and B is indistinguishable in colour from it, then B is C. Let us call this the *maxim of looks*. Now take any two adjacent items in the series. If the maxim is right, then these must be the same colour. For example, *if item  $n$  is red, then item  $n + 1$  is red*. But item 1 *is* red. So, by dispensable use of induction, all the samples are red, which contradicts what we see when we look at the last one. (To get more contradictions, start with the last sample and go through the same procedure in the reverse direction. So the first sample is orange.) Given the maxim of looks, we are *required* to accept the relevant principle, *if  $P(n)$  then  $P(n + 1)$* . So either the maxim is wrong, plausible as it may be, or the above thought about vagueness is on the wrong track entirely.

The maxim may seem even more secure if we apply it, not to 'is red', but rather to 'looks red'. After all, how could something be indistinguishable from something else and still look, or look to be, a different colour? If A looks red, and B looks identical in colour, then surely B looks red. But then, the argument with ensuing contradictions runs as smoothly as before.

This is the moment, perhaps, to inject two thoughts one *might* entertain for at least a second—thoughts which would have stronger appeal in later contexts. Now we can give fairly convincing reasons for dismissing them. But they will need fuller discussion later on.

First, if we concentrate on *looking* red, etc., rather than actually being it, someone might get the idea: why shouldn't something just look red *and* look orange? In that case, we can just accept the maxim *and* the argument *and* the conclusion; what we have is a result and not a paradox. But against this there are two considerations: (1) That *is not* how the last sample looks. Just look! It is reason that is telling us here what the eyes do *not* reveal. (2) Why cannot something *be* red and be orange? Answer: it is a basic part of the notion of red that red is a colour other than orange. So being red excludes being orange. Suppose someone *claimed* (sincerely) that the last sample was red and was orange. Then *ceteris paribus* we would take him as exhibiting his misunderstanding of

shown. I have the impression that Dummett actually thinks there are *valid* versions of the Sorites for 'vague predicates', and thus that, insofar as language contains such things, it is inconsistent. The main purpose of the present work is to show why that is certainly not so.

the notions involved. We would not take him to have described some state of affairs such that one could check to see whether it obtained. But all this carries over *in toto* to looking red. To look red, one would think, is *inter alia* to have a certain look. But further, it is to look a certain colour: namely, red—a colour other than orange. To look orange is also to have a certain look, but a different one. So if something looks red (to you) then its look (to you) is something other than orange. Further, anyone who understands what red is must understand that to look red is to look other than orange. So if he describes something as looking red and looking orange, on *his* understanding of the terms, then we cannot understand him as giving a consistent description of the look something has to him. If he does *not* understand what it is to look red, as manifest, perhaps, by his making remarks like the above, then, *ceteris paribus*, when he uses terms like ‘looks red’ and ‘looks orange’, we cannot understand him as describing some identifiable way things might look at all. In short: the reason why something cannot both look red and look orange is that there is no such look to be had by anything.

The above idea *might* be prompted by thoughts of the following kind of case. Suppose you are asked, ‘Is London a nice city?’ If you are like me, you might well have the inclination to reply, ‘Yes and no’. In so doing, you will not have contradicted yourself. But nor will you have asserted that London is a nice city and that it is not. We will have occasion to return to this sort of remark later.

Second, when it comes to the looks of things, one *might* think it worthwhile to inject tense into the picture. So far, we have spoken indifferently about samples looking red or looking orange. But just *when* do they do this? One might think: if you work your way pairwise through the series, in a series of datable viewings, you may come to a sample which, held next to its (red) predecessor and indistinguishable from it, therefore (thereby?) looks red—*then*. Later, viewed alone in another viewing, it may *then* look orange. But there is no contradiction in looking red at one time and looking orange at another. This thought, too, cannot be right. First, we can suppose that no sample ever appears to have *changed* colour. Each always looks to be the colour it always did (and was). So if the last sample looks orange now, then we can conclude that that is how it *looked*, too (there was no funny light). Second, remember the previous remark: if anything, it is an argument, and not the eyes that persuades us that, for example, the last sample ever looked red.

Finally, it is worth asking exactly what the function is of the premiss that adjacent samples are *indistinguishable* to the eye. Suppose they were just barely distinguishable. Wouldn’t there still be the same problem? It would still be equally embarrassing to try to find that first sample which, while a *near* lookalike to its predecessor, was not in fact, or did not *look* red. So there would remain *some* pressure to continue through the whole series, inferring the redness of each sample from that of its predecessor. On the other hand, this is a situation that strikes us as *familiar*. We all have at least a vague belief about colour that, for

example, red just shades off into orange without there being any sharp divide between the two. That is why we (or some) are inclined to appeal to vagueness here. It remains, of course, to explain just what this ‘shading off’ does to block the argument. But I think we all have confidence that such will, eventually, be forthcoming. The present case, though, looks more serious. For if we have already decided that something is red, then questions of whether something indiscriminable from it is also red can neither hinge nor trade on vagueness in the notion of being red—at least not any vagueness caused by ‘shading off’ phenomena. The maxim of looks suggested that there is one way in which such questions are always to be settled. A parallel maxim where shading off was involved—e.g., involving *nearness* in colour—would, I think, strike us as much less plausible. If we accept the maxim, then paradox (at least) seems inevitable. But why *should* we accept it? It is to that question that I turn next.

### 3. Observation

In our present hour of need, one natural thought would be: reject the maxim. For the culprit in the present case is surely the fact that physics is finer than the eye. But if so, why can it not arrange for fine differences which make for differences in the colour something is, which the eyes cannot detect? But such a move stumbles on several objections. First, Dummett, at least, would respond with a challenge: very well, at what point in the series do you want to say this violation of the maxim takes place? A very good question; much like the problem of identifying the first orange sample in the pairwise *nearly* indistinguishable series, except that the question remains equally embarrassing even if we do not insist on identifying the *first* case where the maxim fails. Whatever the colour of the *K*th sample, do you really want to say that the indistinguishable *K* + 1st is a different colour from that?

A second response from both Dummett and Wright would be that predicates like ‘is red’, and if not, then certainly ‘looks red’, are ‘observational’. Wright puts it like this:

Colour predicates, it is plausible to suppose, are in the following sense purely *observational*: if one can tell at all what colour something is, one can tell just by looking at it. The look of an object decides its colour, as the feeling of an object decides its texture or the sound of a note its pitch. (Wright 1975, p. 338)

Wright is not denying, I take it, that one could be wrong about the look of an item—through not looking hard enough, or bad eyes or light. At least he *need* not deny this. But if there are mistakes of this order, they can be corrected by *looking*—by the one who errs, or if he is, for example, colour-blind, then by someone else. What does not make sense on Wright’s view is that something should have the look of a red thing (or simply look red) to an observer where

there is no visual flaw or abnormality at work, and still not be red, or that it could be *shown* not to be by something other than what one sees by looking. Let us try to generalize: if *P* is an observational predicate in Wright's sense, and if *P* is observable through perceptual device *d* (or devices  $d_1, \dots, d_n$ ), then it is inconceivable that an observer, through exercise of *d* ( $d_1, \dots, d_n$ ), should take (come to take) an item *O* to be what *P* says something to be, where no malfunctioning (or abnormality in the relevant media or channels—e.g., lighting) is afoot, while *O* is *not* what *P* would say it to be. Certainly, under such conditions, the presence of features not observable through the relevant devices or media would not be enough to show *P* *not* to be true of *O*.

If 'red', or some relative, is observational in this sense, then it is clear why the above idea of rejecting the maxim is not on. Suppose that some sample *A* is red and some other sample *B* is visually indistinguishable from it in colour. If *B* nevertheless is not red, this must be because of some (visually) non-observable factors—as we were imagining above, physical differences too fine for the eye. Call these factors *F*. You cannot tell by looking whether *F* is present. If it is, then the relevant item is not red. We need not conclude, though I suspect Wright would, that you cannot tell by looking whether something is red. It is enough that there could be something not red such that one could not tell by looking under ideal conditions that it is not red. That is inconsistent with 'is red' being an observational predicate.

Wright and Dummett have various reasons for thinking that 'red', or something much like it, must be observational, and that, in any case, there must be observational predicates in a language. One reason has already been given. Aside from that, the most interesting, I think, concern the learnability of a language. I defer treatment of these for the last section. Despite the initial plausibility of the reasons, however, I think that there *can* be no observational predicates in a language—in Wright's sense.<sup>2</sup> Hence, while I reject the initial idea about *why* the maxim of looks is wrong, I think that it *is* wrong. I will now try to explain. In the explanation, I hope, lies the solution to Sorites paradoxes in general.

For the ideas I am about to expound, I am heavily indebted to Hilary Putnam (1975/1962) and J. L. Austin (1962*b*, lecture X), who some time ago showed in essence why notions of observability like the present one are not going to work. I hope to add little to what they have done. The interest of the present case, though, is that Wright and Dummett provide the means for bringing the point home yet more forcefully. But we can make still better use of the material at hand if we simplify it. One tends to think of Dummett's and Wright's colour-sample

<sup>2</sup> Note that Wright does *not* insist on a sharp observation—inference distinction in the style of a traditional 'logical empiricist'. The belief that *A* is red *might* represent the conclusion of an inference. It *is* part of the notion, though, that, for an observational predicate, there is a certain *primacy* to observation—to looking, in the case of 'red' when it comes to, for example, what being red really comes to. This, I will argue, is what cannot be—at least in the (philosophical) sense in which Wright requires it.

series as quite long. And, as with the series of *nearly* indistinguishable samples, a feeling may arise that if things just go on long enough, magic may be produced—it happens so gradually that we cannot tell at *what* point red slides over into orange; still, at some point it just does. (Which still leaves us facing the Sorites-type argument.) But the *length* of the series depends on the fineness of our powers of visual discrimination *vis-à-vis* the fineness of possible physical differences. That is an empirical matter. For philosophical purposes, it cannot hurt to assume—falsely, perhaps—that our visual powers are extremely coarse. So let the length of the series be 3. There is a first sample, A, a middle sample, B, and a final sample, C. A looks clearly red when viewed on its own; C, clearly orange when viewed on its own. B is indistinguishable from A when held up next to it, but *also* indistinguishable from C when held next to *it*. You can assume what you like about how B looks on its own. Perhaps it seems difficult to classify. Let us now pose this problem: assuming that A is (looks) red, what colour is B? And C? And what colour(s) do they look? The maxim of looks would lead us to the conclusion that B and hence C are and look red. But is this conclusion right?

If the conclusion is wrong, then so is the maxim. But that cannot be if the predicates are observational. The right conclusion to draw from that is what I believe in any case to be so: that the predicates are *not* observational. It remains to show how this can be. Let us begin with ‘is red’. It seems to me that, present problems aside, ‘is red’, just never was a good candidate for an observational predicate. Though *that* point is not crucial to the present case, it may be instructive to consider why not. The main point is that there are various sometimes satisfactory explanations one might give of what being red is. An explanation for a normal human, who just did not know the term, might simply consist in displaying some red things. If this implies that normally one can tell by looking whether something is red, that is in any case true enough. Another sort of explanation might be a scientific one, referring to wavelengths of light, etc. One might feel an inclination to perceive an order here: the first explanation is ‘more central to’ the notion than the second, or ‘part of the notion’, while the second is not, etc.—or vice versa. Be *that* as it may, two important points remain. First, we do, normally, take both explanations to be explanations of the same thing: we do not feel clearly that one is not really about being red at all. Second, either explanation provides us with what may often be adequate reason for taking something to be red: we may often be credited with *knowing* that such-and-such is red on either sort of ground.

Now what should happen if these explanations got into conflict over cases? This spectre, together with the realization that neither explanation is simply dismissible, led some philosophers, such as C. D. Broad, to suggest that ‘red’ is really ambiguous: it has a sense on which the scientific explanation wins over the ostensive demonstration, and one on which the reverse is true. There is an important and correct intuition here. But as it stands, I think we can see why this is a bad account of the matter. Consider a case: suppose there is a local anomaly

in some corner of the universe—Neasden, say. There are things which *look* red to a normal observer, etc., but instruments reveal them not to reflect the right wavelengths of light.<sup>3</sup> Temporarily, at least, people feel themselves in a quandary.

First point: consider some anomalous objects in Neasden. There may be both true and false things to be said of them, in saying ‘. . . is red’, in calling them red. For example, suppose you want to wear a red tie, or dress, or whatever to the Neasden Deb Ball. Someone hands you the article and says, for example, ‘This tie is red’. It *looks* red all right. And it is not just your eyes or the lights. What more do you want? *Ceteris paribus*, I think we would take what was said here to be something whose truth was immune to measurement of wavelength. Such would not bear on what was said *in this speaking*. *But*: Suppose you want an object to calibrate an instrument, or measure film sensitivity, or . . . , where such is not meant just for Neasden purposes, and the object has to be red. Someone might hand you the same tie and say, ‘Here’s something red’. But here the tie just is not the right sort of thing at all. Then, *ceteris paribus*, what was said to be so of the tie in *this speaking* is not so. So, there is more than one thing to be said in saying ‘It’s red’. Recall our initial fourth proposition: for a given predicate, on given meanings of the word in the language, there may be various things to be said in speakings of it, where what is said is determined not only by the meaning of the words, but also by circumstances of a speaking. Given what ‘is red’ means, we know that we are talking about being red. The problem is, what counts as being *that*? The answer: given what it is normally appropriate to say about being red, sometimes one thing (state of affairs), and sometimes another. *That* point does not turn on dramatic conflicts like the Neasden case (see the discussion of ‘weighs 80 kg’).<sup>4</sup>

Second point: as more facts of the Neasden case emerge, we may or may not be able to come to a general decision as to whether the anomalous objects are red or not. Explanations of the discrepancy may show one of the normally appropriate accounts of being red to be a more reasonable account of what being red comes to under these circumstances—*if*, under them, one account is in general more reasonable than the other. What we cannot do is to know in advance of the explanations which account, if either, this will be. Such is not something accomplished for us by knowing what we are speaking of in speaking of being red. Which, I take it, was Putnam’s point.

Note some ways the above militates against an ambiguity view of ‘red’. If ‘red’ is ambiguous, then there can be no conflict, and hence no resolution of it in the Neasden case. There is a sense in which the anomalous objects, without any doubt, are red, and a sense in which they similarly are not. Further explanations

<sup>3</sup> Compare the following case: is violet a shade of purple? It always seemed so to me. But that is not what the physics texts say. If, after I have read the texts, it still looks purple to me, who is right? Sometimes I think we should say one thing about this matter, sometimes another.

<sup>4</sup> For further discussion of these points, see Travis 1981c and Ch. 1 above.

of the anomaly can have no effect. Nor can purposes for or circumstances in which we call something red. If someone says 'It's red', then, at most, the only question to settle in deciding the truth of his words is which of the two senses they bore. 'In a sense it is, and in a sense it isn't' (Broad's reaction) is a natural reaction to the embarrassment we rightly feel when asked now, for no particular purpose, and in advance of any explanation of the oddities, whether one of the Neasden objects is 'really' red or not. But an equally natural reaction might be: 'You *could* say it's red, and you could say it isn't.' It is the second inclination we have followed up above, with what I take to be its main implication: sometimes we *would* correctly say the one, and sometimes, equally correctly, the other.

Above it has been suggested that a closer look at the phenomena supports the second view rather than the first. First, though we cannot here say whether the objects are red, *tout court*, there are occasions where they would be taken/said correctly to be red, and occasions where they would correctly be taken/said not to be, since there are occasions where one set of considerations rather than the other is correctly taken to show what being red comes to, hence what *is* red, and occasions where the reverse is true.<sup>5</sup> This variation could not be unless there were one thing throughout, being red, about which things were being shown. What better thing to indicate with our word 'red'? Second, it is conceivable that a generally correct resolution of the conflict could be forthcoming—one valid on enough occasions to show us, as philosophers in our studies, what to say about whether the Neasden objects are red. But this supposes that there *is* a conflict, which can only be the case if 'red' is not ambiguous in ways like the one Broad suggests.

#### 4. Looking

The above points are probably disputable. But the advantage of our present tools is that we need not dispute over them as such. For there is a more general point to make. Let us shift attention from being red to *looking* red. Surely, one might think, *this* predicate must be observational. And I suppose, it *is* something like analytic that, normally, but with exceptions, looking red is what something can be seen to do—by *looking*. (The double use of 'look' cannot be an accident.) Still, we are well advised to heed two of Austin's admonitions on this point (1979/1946, esp. pp. 90–4).

First, for present purposes, our concern is with *predicates*, or, if you prefer, concepts. It is a triviality that everything looks the way it looks. Generally, though not always, you can be said to *see* how something looks, if you look. (But: 'I couldn't see how her dress looks because it was all crumpled in a heap

<sup>5</sup> Note: here I use and do not mention, a term ('red'). And if, as I hope, I have said something, then I have used it univocally. (For clarity: it is the colour I am talking about.)



on the floor.’) But to engage the problems round our colour series, anomalies, and so forth, it is necessary to *say* how things look. And what is wanted is not the triviality ‘It looks the way it looks’. So we need to find the descriptions, etc., which fit (e.g., what we see). As Austin points out, the eyes do not *say*—and so are not mistaken about—what things look like. The notion of error engages with what we say or think about given cases.

Second, when it comes to what we say or think, the possibility of error *does* enter. ‘It looks red’ may sometimes be a weaker claim than ‘It is red’, but then, weaker in specified respects, not in *every* respect. In particular, Austin points to two general ways we may go wrong in saying how something looks.<sup>6</sup> First, we may choose a predicate or description which does not fit the look in question: ‘You call that *umber*? (say it looks umber?). *That’s* not the way umber looks at all.’ Second, we may get the look of the thing wrong—not necessarily that we think it looks other than it looks, but, more importantly, that we take it to be a look of a sort that fits such-and-such a description, when it is not the right sort of look at all: ‘That doesn’t look like burnt umber. Look again: it doesn’t look a bit burnt.’

If we now consider our colour series ABC, we can see that Austin was right in a very strong sense. To further the discussion, let us assume that A looks red. Now let us ask what colour B looks, and what colour C looks. As for B, if we hold it up against A, we cannot tell the two apart. This gives us a very strong reason for saying that B looks red—just as looking at the tie in Neasden provided us with a strong reason for saying that it was red. If this exhausts the reasons in play, then we would, so far, be justified in saying that B looks red. But now suppose that we hold B up next to C. Again, we cannot tell the two apart. C strikes us as clearly orange; with respect to its clearly being some colour, it is fully on a par with A. So we now have an equally strong reason to say that B looks orange. We now see that, with respect to what colour B looks, there are two relevant and *prima facie* contradictory considerations (contradictory if, as we do suppose, looking red entails not looking orange). What happens when we hold B next to A gives us reason to say B looks red; what happens when we hold it next to C gives us reason to say B does *not* look red.

Under these circumstances, we cannot be correct in saying that B looks red (and what we thus say will not be true) unless, under the circumstances, we are correct in resolving the dispute in favour of the first consideration and against the second—in holding that putting B next to A shows what is really going on, while putting it next to C does not—or unless it is a determinate fact that the right conclusion to draw is that, counter to our intuitions, looking red does *not* exclude looking orange. Holding B next to A, we may feel a strong *inclination* to say that B looks red. But this inclination will be wrong unless it happens to correspond to the right way of weighing up reasons for and against doing so.

<sup>6</sup> On these ways of going wrong, it is also worth consulting, Austin 1979/1952.

There are uses and uses of 'looks red'. For some, surrounding circumstances may point to a right way of weighing up the reasons—and different circumstances may point to different ways. So there may be true things to be said in saying B to look red, in *some* speakings of 'B looks red', as well as false things to be said in so saying in other speakings. Particular circumstances aside, neither reason counts as any stronger than the other. But then the conditions are not fulfilled for us to be able to say, correctly, that there is a fact as to whether B looks red. In that case, we must say what there is to be said: it is not a fact or true that B looks red, or a fact nor true that it does not. Such *might* be expressed in saying 'It looks red and it doesn't', not by way of asserting a contradiction, but as a colloquial means of withholding judgement—as recognition of the fact that B neither counts as fitting (satisfying) the predicate nor as failing to.

We can also, perhaps, do justice now to the intuition that somehow tense should get into the picture. Our original predicates, 'is red' and 'looks red', under the above circumstances, leave us at a loss for words. It is natural to search around for others. Some fairly ready to hand (or mind) would be: 'looks red when you hold it next to A' and 'looks orange when you hold it next to C'. The first is naturally understood so that observations of B next to C are irrelevant to what it says to be so. So understood, it is true of B—since a resolution of our conflict is built into what it means. Similarly for the second predicate. Here the occurrence of 'when' suggests something like a reference to time. But it cannot really be *time* that is involved, as evidenced, first, by the fact that we are not supposing that B changes the way it looks over time, and second, by the fact that it is still true that B looks red when you hold it next to A, even at a time when you hold it next to C: such is something it does all the time; a fact about how it looks *tout court*. And that, I think, is about as far as time takes us.

Let us now consider C. Here, again, there are reasons for and against taking it to look red, or taking it to look orange. Here, however, there seems to be a clear and general resolution to the conflict. The reason for saying it looks orange is just how we take it to look when we look, no matter what it is held against. The reason against so saying is an argument, rehearsed above, making appeal to the maxim of looks, and through that to the observationality of 'looks red' and 'looks orange'. Clearly, here it is the argument that has to go. Except perhaps in the most exceptional circumstances, it is true to say that C looks orange, not true to say it looks red, full stop. From which we can conclude that the maxim is wrong, hence 'looks red' is not observational. (Which does not mean you cannot generally tell the looks of a thing by looking.)

The above treatment may seem *ad hoc*. Given that there are reasons for saying B looks red and reasons for saying it does not, why not heed both? One might even say for each reason that it is built into the meaning of 'looks red' that it *must* be heeded, come what may. The conclusion would be that 'looks red', etc., are contradictory predicates (notions). But we need only ask ourselves what we *know* about what 'looks red' means. Does it mean something such as to apply

to B? It *looks* as if principles which we elsewhere appear to accept dictate that it does. So this *seems* required if we want to treat like cases alike, *if* we attend to the comparison with A. But similarly, it looks as if principles we elsewhere seem to accept dictate that B does *not* look red—if we attend to the comparison with C. If we embrace both appearances and maintain conclusions, then it is not just the predicate or concept which is contradictory; *we* have contradicted ourselves.

The way to avoid committing this mistake is: don't do it. Neither language nor logic compels it. We *appear* to have reasons for saying each of two contradictory things. But then, we could take either or both of these phenomena as mere appearance. Not doing so lands us in contradiction. What better, more powerful reason could there be for doing so? It only remains to find the most reasonable way of sorting out mere appearance from fact.

But then how to describe B? For many purposes, you might say: you cannot exactly say it is/looks red, and you cannot exactly say that it isn't/doesn't—a point often to be put as: it is (looks) red and it isn't (doesn't). For many, though not all purposes, this is enough. Where not, we can always repeat the details of the case. But, in sum, languages, like eyes, do not speak; *a fortiori* not contradictions.

If we accept the above, then 'looks red' cannot be observational in Wright's sense. For how do we tell that it is not a fact (or true to say) that B looks red? Not by something we can notice when we look at B, but despite anything we can notice, by attending to an *argument*—one concerning how the pros and cons of applying 'looks red' in this case are to be added up. The way to tell it is not true that B looks red is *not* to look harder. But arguments, where they exist, *must* be attended to. A good reason for not applying 'looks red' to B is a good reason—not an ignorable one when it comes to whether B looks red. That, in essence, is why there can be no observational predicates in Wright's sense.

Finally, if all the above is so, then is 'looks red' a vague predicate? As a bit of English, I think, it is neither vague nor precise. Nor is it clear what a *predicate's* being the one thing or the other should come to. *Of course*, as with any predicate, there are some imaginable cases where what it means does not make it a fact that it applies to some item, or a fact that it does not. As Wittgenstein (1953, §84) wondered, what would a predicate look like for which *that* was not so? If we turn our attentions to speakings, we have described three cases, and there is a fourth. First, someone might say 'B looks red', where what was said, on a proper understanding, is such that the comparison with A, but *not* that with C, is to be attended to in evaluating whether things are as said to be. In that case, if our interests are confined to evaluating truth, then what was said is perfectly clear; not a bit vague. Second, for some other speaking, what was said, on a proper understanding, may be such that the reverse is true. Then the same comments apply. Third, for another speaking, it may be clear that a proper understanding is such that neither comparison outweighs the other. Then it is clear of what was said that it is not a fact that it is true, or a fact that it is false. So far, no ground for talk of vagueness. But fourth, for some other speaking,

it might be unclear which of the above three descriptions it fits. In which case we *might* say the speaker was vague as to what B looked like. Further, in any of the four cases, a speaker *might* justly be accused of vagueness if he was not clearly stating what we wanted to know. (We asked whether it was red; he said that it *looked* red.) So one could speak vaguely in saying 'looks red', but one need not.

## 5. Phenomena

There may be a residual feeling: perhaps 'looks red' is not observational. But then surely something else must be. 'Seems to look red', perhaps? A good way to dispel such thoughts, I think, is to try to *introduce* observational predicates into the language, just in case there are none yet there. Let us, then, introduce the terms 'phred' and 'phorange' ('ph' for *phenomenal*). Now it is up to us to *stipulate* what they are to mean. How should we do this? Since we want these predicates to be observational, we can try just stipulating that. So, by stipulation, you can tell by looking whether 'phred' and 'phorange' apply, and nothing other than looking can controvert what your looking reveals. (One is almost inclined to say: if you take it to be phred, then it is. But, of course, that only raises problems about when you are taking it to be *that*.)

But obviously the above is not enough. If experience is to decide where 'phred' applies, we need still to say something about what sorts of (or which) experiences these are to be. Intuitively, as orthography suggests, 'phred' is to have something to do with red, and 'phorange' something to do with orange. But exactly what? For a start, like 'red' and 'orange', 'looks red' and 'looks orange', 'phred' and 'phorange' are to be mutually exclusive descriptions. But that is not enough to say where they fit. There are two courses open: we can try to explain this, or we can try to show it. In the first case, we might say, 'Phred is what something is when it gives you the inclination to say it looks red.' (Don't worry about inexactness here; nothing will turn on a nuance.) In the second, we will simply exhibit some things which *are* phred. The first explanation may leave us with problems, in certain cases, as to just what your inclinations were. The second leaves us with a potential problem, for novel items, as to whether they are what was being exhibited. Still, if you know a better way of laying down stipulations, feel free to try it out.

Now let us confront our stipulation with the series ABC. Whatever explanations we attend to, if they have forged the desired link with red, we will be inclined to say that A is phred. But now our stipulation of observability appears to take over. B, since indistinguishable from A, is also phred. By like reasoning on B, so is C. But the process is reversible: C, if our explanations worked, is phorange. Then so is B, and so is A. So we are back in our predicament, but this time with different predicates.

What the predicament shows is that we must be misunderstanding what the stipulations tell us about this case. Again, let us concentrate on B. Holding it up against A gives us excellent reason to take it that B is phred. But holding it up against C gives us equally excellent reason to say B is not phred—in fact, that it is phorange. So now look at the explanation by showing. Is B a case of what was being exhibited in those samples—is it the same thing as they were? The first consideration gives us reason to answer yes, but the second gives us reason to answer no. Similarly for other sorts of explanations. Is it really phred that we are inclined to call B? There is a reason for saying this and a reason against. Unless one of our reasons can be seen to outweigh the other—as, *in abstracto*, it cannot—we cannot take it that B is phred, or that B is not phred; B neither counts as being nor as not being that. Similarly for phorange. When we get to C, the reasons in play are thus different. But the reasons for calling it phorange seem to outweigh anything else. So, in this case, ‘phred’ and ‘phorange’ turn out to behave just like ‘looks red’ and ‘looks orange’.

Just a minute, however. As argued previously, if we allow ourselves to weigh up reasons in the above way, then we are not treating ‘phred’ as an observational predicate. So doesn’t our stipulation that ‘phred’ is observational forbid us from doing just that—in effect, from listening to reason? My own inclination here would be to take it that this was a misunderstanding of what was meant in the stipulation by ‘observational’. It could not have been the intention that we were required to behave irrationally. Probably the stipulator just did not have this case in mind.<sup>7</sup> But suppose you reject this line. Then it is mandatory for you simply to obey your inclination to call things phred or not no matter what. Where phred is concerned, there is for you no difference between seeming right and being right, if I may borrow an idea. So, though you may sometimes say ‘It’s phred’, or ‘It’s not’, we can no longer assign any content to your noises (i.e., there is none). Sometimes, for example, when B is held next to A, you may say ‘It’s phred’. Other times, you may say ‘It’s not phred’. But by stipulation we are forbidden from engaging in the reasoning necessary for determining whether you have or have not contradicted yourself. That, of course, depends on just *what* you said each time, about which, by stipulation, there can be no reasoning. Nor are we entitled to attend to the argument, modelled on that for ‘is red’ or ‘looks red’ which shows that there might be some problem about whether C is phorange. We are no longer in an arena where *problems* could arise at all.

So why can’t there be observational predicates (in which things may be said to be so)? Because it cannot be stipulated, or otherwise guaranteed, that difficult cases, like Wright’s and Dummett’s colour series, where conflicting reasons for judging come into play, cannot arise. Where they do, our concepts, or more accurately our judgement, is prepared. We cannot be obliged, even by stipulation,

<sup>7</sup> Don’t be worried by the fact that the stipulator was, apparently, *me*. That changes nothing.

not to exercise good judgement or good sense. Apparently faced with such a demand, it is we who are responsible for keeping our wits about us.

## 6. Sorites

What is there in the above to resolve Sorites puzzles? First, why doesn't the problem arise for the three-sample series? Since it makes no difference, let us from now on talk simply about *being* red or orange. Originally, then, we were asked to assent to a principle which, given the undisputed fact that A is red, would force us to conclude that B is red. Specifically, the principle was: if B, held next to A, is visually indistinguishable in colour from it, then B is the same colour as A. On examination, however, we concluded the following: there are good prima facie reasons to accept the principle. Hence, the principle (plus the observation) gives us good reason to take it that B is red. *But* we also have, from another source, equally good reason (*in abstracto*, at least), for taking it that B is orange, hence not red. Therefore, despite original appearances, we are not correct (again, *in abstracto*) in taking it that B is red. So, we do not quite accept the prima facie plausible principle either. Now we come to the move from B to C. Since we do not take it that B is red, we do not have *that* reason (plus indistinguishability) for taking it that C is red. Still, we have *some* reason to take it that B is red, and this may be seen as transmitting to C some reason to take it that C is red. The reasons, such as they are, though, are clearly outweighed by our reasons for taking C to be orange. And (hence) there are no good reasons for taking C to be both.

Thus we reject any principle, no matter how plausible, which, together with undisputed facts about A, would compel us in general to assert that B is red. But, as for the above principle, the above thoughts suggest something to put in its place: *not* 'If A is red, and B visually indistinguishable from A, then B is red', but rather, 'If A is red and B visually indistinguishable from A, then that provides excellent reason for judging that B is red'. As long as we stick to reasons for judging, and concern ourselves with how judging is to be done, paradox simply cannot arise.

The above thoughts are not affected by making the series longer. Suppose, with Dummett and Wright, that there are many colour samples. Now consider any indistinguishable pair. If one member is colour C, that gives us excellent reason to take it that the other one is C. What effect does this have on the whole series? Again we suppose that the first member is unmistakably red. That, plus indistinguishability, gives us excellent reason to say the same of the next member. Such is transmitted not much diminished to the next member, and so on. As we go through the series, however, the force of the reason transmitted by this process diminishes. The reason is that a different, though parallel chain of reasons is transmitted from the last member of the series, which is indisputably

orange, in the opposite direction. So, for any member of the series, there will be some reason, no matter how slight, for calling it red, and similarly for calling it orange. Since the series is long, there may be many cases where one reason clearly outweighs the other. The first and last members are two such. But there may also be many cases where weighing yields no clear result.

We can now see just where thoughts about vagueness might enter. Weighing up reasons pro and con is generally not a precise business. There seems no question of metricizing strengths, at least in cases like the above. Further, despite the fact that we do not often disagree in our weighings-up, when it comes to judging being red, for example, we may feel some subjective element in our perception of weights. Last but not least, for a given item, and given reasons for and against applying a predicate to it, we may weigh these reasons differently under different circumstances of weighing, or for different speakings. There is nothing in these facts to suggest that the *predicate* is vague. For any predicate, there could be reasons for and against applying it in some given case. Whenever reasons must be weighed up, the intrinsic impreciseness of that task comes ineluctably into play. The result may be that there are many items—e.g., in the colour series—such that it is unclear, or unclear for some given purposes, how the reasons do add up with respect to them—all the more so if such *is* clear for other purposes, but then with different outcomes for different purposes. So of course we may be embarrassed if asked to pick out that item in the series for which reasons cease to add up one way and begin to add up the other. But there should be no surprise in that.

Appeal to indistinguishability of colour really represents a special case. Looking at the long colour series again, let us generalize it away. What is true of adjacent samples in this series is simply that, on general grounds, there is good reason to think, for a given predicate, that if the predicate applies to one sample, then that is good reason to take it to apply to the other.<sup>8</sup> The same would apply if the samples were not literally indistinguishable, but merely very very close—for example.

With the generalization, we are now prepared for Wang's paradox. This is an argument with two premisses and appeal to induction. The premisses are: (A) 1 is a small number. (B) If  $n$  is a small number, then  $n + 1$  is a small number (for all  $n$ ). The first thing to note is that this paradox derives much of its plausibility from the fact that, like most philosophical conundrums, it is stated outside of any context of use. What about (A), for example: is 1 a small number? Consider the following: one air traffic controller to another: 'How many jumbo jets have crashed today?'; other: 'Only one'; first: 'Well, that's a small number'.

Seeing whether a number is small *tout court* is more than just a difficult task, and troubles begin with 1. The problem is not that matters get less clear by

<sup>8</sup> Unnecessary duplication of 'good reason', perhaps, but I want to emphasize where the elbow room comes.

degrees as numbers get bigger. For there is no degree of clarity or justness with which 'is small' fits a given number as such. Rather, there are many ways of marking off the small from the non-small, each correct for some judging of the matter. What is unclear on one way may be quite clear on another, and what is clearly small on one way clearly not on another. The multiplicity of ways such matters are sometimes to be decided naturally makes it hard to spot a unique boundary which is correct *tout court*—not because it is unclear where such lies, but because there isn't any. Where circumstance indicates a particular way of bounding the small, there may still be unclear cases, of course. But at least there may also be clear ones.

With the above reservation, let us let (A) pass and turn to (B). In view of our general discussion, (B) really ought to be rephrased: since, for most purposes, 1 does not make a great difference to the size of a number, if we take it, for some  $n$ , that  $n$  is small, we thereby have excellent reason to take it that  $n + 1$  is small. Such reason, of course, is subject to coming into the balance with other reasons, including reasons against calling  $n + 1$  small. If there are such, it will have to be weighed up with them. What the other reasons are will depend heavily on context and purposes. Which is why it would be nice to have some on hand. Suppose, for example, we are interested in the number of people who inhabit various cities. Then for many purposes, we may be able to agree that 7,000,000 is a large number. This means that any reason for calling 7,000,000 small, derived from its proximity to small numbers, or numbers which there is some reason to call small, is clearly outweighed by other considerations. If 7,000,000 is large, then it transmits reasons to its neighbours downwards by the same considerations by which 1, if it is small, transmits reasons upwards. The number 6,999,999, for example, is quite close to 7,000,000, which provides excellent reason for saying that it, too, is not small. For any given number between 1 and 7,000,000, the competing reasons from these two sources, and any other reasons, if they are present, will have to be weighed up. Even though it is numbers we are concerned with, adding up reasons is not an arithmetical process, or a precise one. So there may be many cases where we cannot perceive any clear outcome of the weighing. That does not matter, so long as there are cases where the outcome is clear, such as 2 and 6,999,999. The principles which it is reasonable to accept (not quite premiss (B)) can thus be seen not to entail that all numbers are small. There is no paradox.<sup>9</sup>

At the outset, we mentioned two ways the meaning of a predicate might underdetermine what it is true of. First, the predicate may be 'vague', on the conception of vagueness attributed to Dummett. I have suggested that every predicate is that. Second, where the predicate speaks of being  $\Phi$ —e.g., being

<sup>9</sup> The point remains if there are no purposes at hand, with the difference that in that case, for most numbers, if not all, we will have no idea how reasons should be toted up. That does not mean that we *should* have the idea that, really, all numbers are small.



small—there may be various ways of viewing what being *that* comes to, each of which is, on occasion, the right one for judging what is and is not  $\Phi$ . This too, I think, is true of every predicate. For dealing with the Sorites in its various forms, such as Wang's, the point is not only true but important. Give us reason to speak of smallness in one particular way, among the many sometimes correct ones, and we thereby also gain reasons for taking some numbers not to be small. Where these must be weighed in the balance against any generated by revised (B), Wang's argument (with revised (B)) no longer drives us inexorably to paradox. Where we have no reason to speak of smallness in any particular way, we also have, *ipso facto*, no reason to take reasons for calling a number small, insofar as these are generated by revised (B), as decisive. No way of weighing up reasons is just that. So, with no way of speaking about smallness, there is nothing to be said about which numbers are small—*inter alia*, nothing paradoxical.

The position I hope now to have reached may be summed up as follows. It may *look* as if languages contain contradictory instructions for their use. But what this comes to is that what a particular predicate means, or, better, what we know about this in understanding it, may provide us with reasons both for and against applying it to some particular item. This is not specifically a property of 'vague' predicates, but may be a property of any predicate whatever, if the world furnishes us with the right items. Where reasons conflict, we must have some way of weighing them up. In general, there may be many such, each of which sometimes would be proper. One such *may* be indicated by the circumstances of a given weighing. Given it, some reasons may be seen to outweigh others, hence to be decisive for the applicability of given words to a given item. Where such is to be seen, what we may see thereby is whether what would then be said of the item in speaking the predicate, on a given understanding of it, would be true, accurate, or just, and with what degree of clarity such things would be so. When we try to abstract such judgements from the circumstances which show how they are to be made, we arrive only at ill-defined notions, e.g., of clarity. Of course, paradox will result. But the paradox then lies in our methods of judging, not in the language.

## 7. Learnability

As mentioned earlier, Wright offers an explanation of why there *must* be observational predicates. We would be too swift in concluding that there cannot be if we did not say something about why the explanation does not work. Wright's case consists of two points. First, if language is to connect with, or be about the world at all, then there must be some predicates which are learned through experience—i.e., by being shown cases where they apply or do not. Whatever the merits of this idea, it will not be disputed here. Second, showing you situations where a predicate applies presupposes that you can see, by observing them,

whether these are situations appropriate for the term to apply or not. Wright puts it like this: 'It would be a poor joke on the recipient of an ostensive definition if the defined expression applies selectively among situations indistinguishable from one which was originally displayed to him as a paradigm' (Wright 1975, p. 341).

There are here several points to dispute. For example, the idea that one cannot 'connect language with the world' without learning some predicates through experience does not entail that there are some predicates (in a language) such that one cannot 'connect the language with the world' without learning *them* through experience. So the need for predicates in a language which, as such, just are learned through experience is not obvious.

A more significant sign that something is going wrong is the following observation. Suppose we compare 'is red' with 'looks red'. If we had to choose just one, most of us, I think, would pick 'looks red' as a more likely candidate for an observational predicate than 'is red'. If we think about how these predicates are learned, however, it is a plausible though not *a priori* deducible hypothesis that 'is red' is learned first, then 'looks red': first, the bold and direct, though perhaps naive characterization; later the cagier hedging. Things *could* be like that, which shows that it might be a mistake to connect observability of a predicate with how it is learned at all.

The main issue, however, is epistemological. Why would Wright's 'joke' be a poor one? I have no idea how language learning goes, except that I do not think it is a philosophical matter. But why couldn't it be like *this* for 'is red'? The learner is shown samples of things which are red. The samples—and the situations—are carefully chosen so that in these cases, you can tell by looking whether the items are red. And it is what you see when you look that tells you so. If you went on to say you knew the items were red, you would not be wrong (*ceteris paribus*). What the learner *learns*, however, is this: these are cases of red objects (i.e., where 'is red' applies). What *makes* these red objects is what you see when you look. Of course, there might be other considerations which count for or against something being red. These the learner may hope to find out about in time. What he might have to do to do this is, on occasion, when presented with a candidate for being red, to pose the question, of certain other considerations: 'Are these reasonably taken to count for or against being the same thing as what I formerly observed counted for or against being?' Occasions, as well as answers, are likely to suggest themselves from time to time. Which is to say that the meaning the learner assigns to 'is red', even if he learns it ostensively, through experience, need not be exhausted by what he observed in the cases he was shown.

The main thing *obviously* wrong with this story is that language learning, we know, is rarely and need never be explicit. The meaning a learner actually assigns to 'is red' must, it would seem, be distilled out of explanations of all sorts (including wrong or suspect ones, garnered catch-as-catch-can from all possible sources). So much the less reason to think 'is red' an observational predicate.

If what the learner learns is as described, it is conceivable that there could be items indistinguishable (by looking) from the ones he was shown where 'is red' does not apply. Where is the joke in this? *If* we showed the learner situations where he ought not conclude, on the basis of what he sees, that the item in question is red, and if we did this in such a way that he took it he could conclude this, then later, perhaps, we could all have a good laugh. But such need not be the case at all. Let us end on an Austinian note. I want to teach someone what a pig is. So I show him some. Later, Dr Scherzenstein builds a very clever mechanical pig. When you see it coming around the corner of the barn, you cannot tell it from the real thing. Knowing Dr Scherzenstein, one ought to be cautious. In his barnyard, you might not be able to tell by looking whether something is a pig. But need our learner suppose otherwise? Other cases are other cases, and we need not have blinded the learner to that. Nothing in our training properly taken would suggest that he should so suppose. And properly taken, here, is one way humans normally do take such training and examples. The Scherzenstein case does not show that one cannot, sometimes, just see that something is a pig. Nor does it suggest that in the original cases, from which the meaning of 'pig' was learned, one ought to have doubted whether the things exhibited were pigs, or that what the learner saw did not show them to be pigs. There *could* be other relevant considerations. In the Scherzenstein case there *are*. But why should our learner, who knows what 'pig' means, be barred from finding this out? Knowing what 'pig' means does not mean knowing, for any case whatever, what would show whether something was one. It had *better* not mean this if *we* are to know what 'pig' means. The suggestion is that in these respects, being red and being a pig come to much the same thing.

# 10

## Attitudes As States

A theme recurs in Wittgenstein's later work:

If one says that knowing the ABC is a state of mind, one is thinking of a state of a mental apparatus (perhaps of the brain) by means of which we explain the *manifestations* of that knowledge. Such a state is called a disposition. But there are objections to speaking of a state of mind here, inasmuch as there ought to be two different criteria for such a state: a knowledge of the construction of the apparatus, quite apart from what it does. (*PI* §149)<sup>1</sup>

Try not to think of understanding as a 'mental process' at all.—For that is the expression which confuses you. . . .

In the sense in which there are processes (including mental processes) which are characteristic of understanding, understanding is not a mental process. (*PI* §154)

How does the philosophical problem about mental processes and states and about behaviourism arise?—The first step is the one that altogether escapes notice. We talk of processes and states and leave their nature undecided. Sometime perhaps we shall know more about them—we think. But that is just what commits us to a particular way of looking at the matter. For we have a definite concept of what it means to learn to know a process better. (The decisive step in the conjuring trick has been made, and it was the very one we thought innocent.) (*PI* §308)

The difficulty becomes insurmountable if you think the sentence 'I believe . . .' states something about the state of my mind. If it were so, then Moore's paradox would have to be reproducible if, instead of saying something about the state of one's own mind, one were making some statement about the state of one's own brain. But the point is that no assertion about the state of my (or anyone's) brain is equivalent to the assertion which I believe—for example, 'He will come'. (*RPP* §501)

Its subjects are meaning, understanding, and attitudes such as expecting, believing, hoping, and intending. The point is not that there is no good sense in which such things are mental states (*vide PI* §572). It is rather that so thinking of them suggests false pictures. Which ones? At least two. This paper will explore just one, and only for belief.

Wittgenstein's point seems mysticism *if* we suppose a certain sort of scientific discovery may be made. There are a motley of cases in which someone may be

<sup>1</sup> Throughout this essay, *PI* stands for Wittgenstein 1953, and *RPP* for Wittgenstein 1980, vol. 1.

observed to believe such-and-such; we ascribe belief on a motley of grounds. The thought is that there *may* be an underlying phenomenon to be discovered; one the motley all manifest. Believing such-and-such might then be that phenomenon. Denying that there is one then seems to suggest, mystically, that manifestations happen for no reason. But the point does not apply with this thought in place. It is precisely that there is no such open question; what belief ascriptions *say* leaves no room for such discovery. The point is one of proper grammar; good grammar is not mysticism.

Two mutually supporting pictures, though, promote illusion. One is Fregean; the other dates at least from Descartes, and is prominent in the empiricists. Both remain alive and well. The Fregean picture concerns semantic items; items which, while mental or physical vehicles may bear them, are neither mental nor physical. Rather, their identities are fixed by relations such as entailment which they bear to one another, and by their bearing, as they do, the semantic properties they do. Calling these items non-mental emphasizes that each must have an identity independent of the vicissitudes of any given mental life: its being the item it is cannot require that so-and-so takes *this* attitude towards it, or treats it thus. The Cartesian picture concerns items which are mental, and may be physical as well. Its central idea is that a belief requires a medium to represent what it is about.

Each picture proposes a constant in someone's believing such-and-such. It purports to identify something fixed about the believer in his holding that belief. So each offers a state for belief to be. Such a state's nature *might* repay empirical investigation. Which would make Wittgenstein seem a mystic. Conversely, what might thus seem mystical in Wittgenstein targets both views. Here, though, only the Fregean will be examined.

## 1. The Fregean Picture

Frege is thought by some to lack a picture of the mind. Let us make do with what we find. He does say that if someone ascribes a belief—say, in saying 'Pia believes that Esau is hairy'—the words following the 'that' refer to a sense (in fact, their 'normal' one). The grammatical object of 'believe' thus refers to an object. Not every grammatical object makes that contribution to what its whole says.<sup>2</sup> This one does. To be as said to be when said to believe something, then, is to bear a particular relation to a specific item, always of one given sort. The relation is the one the verb 'believe' expresses. To bracket problems, I will call it *relating believably to*. For Frege, there is a fixed domain of senses, hence of objects of the relevant sort. To believe anything is to relate believably to one such object.

<sup>2</sup> E.g., 'The dragon belched fire', 'They had a good time', or 'The hole caused the boat to sink'.

*Senses are semantic items.* Call a set of semantic properties a semantics.<sup>3</sup> Call all of an item's semantic properties *its semantics*. A semantic item is what has a semantics. There are two conceptions of this. On one, the item *has* semantic properties. It is true, say, or speaks of being white. This captures the intuition that if Pia believes Esau is hairy, but he isn't, then her belief is false. On the other, the item *is* a collection of semantic properties. Which captures for senses the intuition that for words to have a given sense just is for them to have the semantics they do. Not choosing, I will speak of bearing a semantics. I distinguish a semantic item from a mere semantics by this: if a semantic item bears a semantic property P, and if having P requires having Q, then the item bears Q; if having P excludes having Q, then the item lacks Q.<sup>4</sup> Given the wall, saying it to be white may require being true. Semantic items, one might say, are closed under requirement.

On Frege's conception, for each sense, a (perhaps not proper) part of its semantics identifies it as being the sense it is; a sense could bear *that* semantics only by being that sense. Such provides *the* correct way of counting senses: different identifying semantics, different senses. Such semantics is meant to be a *fixed* feature of a sense; one it counts on every occasion as having. So, too, for objects of belief.

*Senses are fine grained.* Many semantics contain the properties of referring to, or speaking of, Esau and of being hairy, and of saying that person to be that. Each of these semantics fixes a different sense. Conversely, many senses fix just that set of referents related in just that way, or, as I will say, just those topics. Each such sense constitutes a different object of belief. So there are many items one might believably relate to in believing Esau to be hairy; many distinct beliefs one might thus hold. To relate believably to one such is not so to relate to any other; so to relate to two such is, if possible at all, to hold two distinct beliefs. To say Pia to believe that Esau is hairy is always, if it is to say anything, to say her to relate believably to just *one* specified such item. One who relates believably to some such does not so relate to all; otherwise, senses were uncalled for here.<sup>5</sup>

*Duration.* 'Believe' is a state verb. If belief is what it speaks of, then belief always has at least some of: an onset, a duration, and an end. Onset and end may

<sup>3</sup> What needs to be said about semantic properties will be said in section 4.

<sup>4</sup> If beliefs are items, they are thus semantic ones. If you say N to believe that 'A', but your 'A' is incoherent—without a semantics that something could have—then you have not furnished anything to be believed or not, so have not said N to believe anything. The fault is with you, not N. N may be confused, but is prevented by the concept of belief from holding incoherent ones. There are no such feats to be performed.

<sup>5</sup> The 'evidence' that belief needs objects such as senses consists in this: we can speak falsehood in saying 'Pia believes W', even though our W identify topics to which Pia somehow relates believably (Esau and his being hairy, say). To ascribe to Pia belief in the wrong item, but with the right topics, is, on Frege's view, flatly to speak falsehood. I have explained why this evidence is unconvincing, and the conclusion implausible, in my 'Are Belief Ascriptions Opaque?', Ch. 8 above.

be gradual or absent. The following shows the idea. Hugo, walking in Malet Street, sees a woman he takes for Pia disappear around a corner. He concludes that Pia is in London, and not, as she said, in Paris. Hugo persists in this belief into mid-afternoon. At which time he receives a call from Paris from Pia. So, he concludes, that woman was a ringer; Pia is in Paris after all. From his stroll until that moment, Hugo believed that Pia was in London, a belief he came to have on seeing the woman turn the corner, and ceased to hold on hearing Pia's static-punctuated voice.

*The picture.* The Fregean picture specifies what is fixed about one, for the duration, in believing such-and-such. There are four central points.

*First*, to believe something is to relate believably to some one object of one given sort. That it is *that* object is, or fixes, what is thus fixed about one for the duration.

*Second*, that object is a semantic one. It bears a semantics which identifies it as being the object it is. That semantics is something it bears *tout court*, so a fixed feature of that belief for the duration. On any occasion, any object is that one just in case it has that semantics. So for anything one might believe, there is one particular semantics that constantly identifies *what* is thus believed. Further, there is just one right way of counting objects of belief: count identifying semantics.

*Third*, objects of belief are fine-grained. Their identifying semantics encompasses, and needs to, much more than just the properties of being about their *topics*, such as Esau and being hairy.

*Fourth*, words, saying someone to believe something, actually refer to (mention) some such item, and *say* the person to relate believably to it. Whether they are true thus turns on whether that item is as said to be—believably related to by the subject—so whether that item does for that subject what an item so related to must do. That item plays the designated role for the subject only if its identifying semantics does. So the subject must relate properly to that. Part of that relation is this: that semantics is, *tout court*, the one that identifies the content of the subject's belief, *what* it is he thus believes. Such words thus say their subject to believe what that semantics identifies *tout court*.

If words say N to believe A, then saying what they did requires referring to the very object of belief they did. The possible ways of counting words as doing that are exhausted by the ways of counting items as being that object or not. There is just one way: such items are counted by their fixed identifying semantics. There are limited ways of referring to a given such object: the one given words refer to is deducible from the relation their other semantics must bear to the identifying semantics of their referent. What words 'N believes A' say, then, identifies precisely what semantics words must have to say N to believe that, so what else would do so. Their truth requires that precisely the words thus identified *are* what would say of N what they did.

Senses were made to fit words, for which they were assigned many functions. It is a nice question whether what performs those functions is also fit to be an object of belief. The above, abstracting from most of those functions, does not demand an answer. It fits many views of *which* objects are believed. Some dissent from the third strand, preferring what have come to be known as Russellian objects of belief. That strand has a rough ride ahead. But the present criticism survives without it. Before critique, though, I will expand on what Frege's view implies.

## 2. Beliefs and Expressions

Pia's belief that Esau is hairy may be shaken by fresh evidence. Or grow stronger. Or gradually turn to doubt. Such is one nominalization of 'believe'. Its close relative is the noun 'believing': Pia's believing as she does what she does. The belief that Esau liked porridge was once widely held; but experts are now sceptical. This suggests a second nominalization on which 'belief' means that which is, or may or could be, believed. If 'belief' can mean that, then these are two different nominalizations. The belief that Esau is hairy did not grow stronger, though Pia's belief did. Call the last of these *belief*<sub>1</sub> and the first *belief*<sub>2</sub>.

We sometimes count beliefs<sub>2</sub>: the belief that Esau is hairy is one thing; the belief that he likes porridge is another. Which suggests a way of counting beliefs<sub>1</sub>. To believe that Esau is hairy is not to believe that he likes porridge; so Pia's belief<sub>1</sub> that he is hairy and her belief<sub>1</sub> that he likes porridge are two. Which does not suggest just one right or privileged way of counting either beliefs<sub>2</sub> or beliefs<sub>1</sub>. Frege's picture identifies beliefs<sub>2</sub> as items of a particular sort. Each such item is identified by a particular semantics. There is, then, a unique right way of counting items of the appropriate sort, so beliefs<sub>2</sub>; which suggests at least a privileged way of counting beliefs<sub>1</sub>. The definiteness of beliefs<sub>2</sub> (on Frege's picture) imposes itself naturally on beliefs<sub>1</sub>. Conversely, the exact nature of Pia's belief<sub>1</sub> that Esau is hairy depends on fine points of her psychology, *inter alia*, what she is prepared to recognize as Esau, and how or when she would do so. There is a temptation to read such fine detail into her corresponding belief<sub>2</sub>, which misleads philosophers into supposing that objects of belief must be fine-grained. (Though the natural view is that a belief<sub>2</sub> is what many beliefs<sub>1</sub> might share.)

There are also two ways of speaking of expression of belief. On an occasion, Pia *could* say many things. She could use many words, with many semantics, to do so. In saying some of these, in speaking some of these words, she would express (say) what she believed, so, on one use of 'express', express her belief. In saying/speaking others, she would not. On one use of 'express', she thus expresses her belief only if she believes what she says, and that is appropriately non-accidental. In expressing her belief, she expresses *herself*, which she alone can do. Nor would she have expressed herself if she inadvertently said what she just happens to believe.



There is, I think, a wider use of 'express'. Suppose Dan says 'Esau is hairy'. On one use, it does not follow that he expressed a belief; that depends on whether he believes what he said. He did at least express a thought. That is, he said something; some may believe it, some may not. If Pia is among the believers, then Dan said what she believes. He might then be said, on this wider use of 'express', to have expressed what, or something, she believes.<sup>6</sup> On this wider use, we often express not only our own beliefs, but those of others. Or so we suppose. Doing so requires only that those others believe what we say. Where necessary, I will speak of wide and narrow expressing, which I will write, 'w-express' and 'n-express' respectively.

Pia's belief that Esau is hairy has a duration. Within and without it, many occasions arise for expressing what she then believed. Suppose that, one May, her belief is extinguished by an encounter with the strikingly smooth Esau. The following May both we and she might w-express what she formerly believed about Esau, though neither she nor we could n-express it. Occasions are not times. Many occasions for expressing Pia's belief may arise, or might have arisen, at any given time: Dan saying to Hugo 'Esau is hairy' in one situation, while Sam says the same to Eva in another. The time for such an occasion need not be one when Pia believed that.

Expressions of belief are distinct from ascriptions of it: words *saying* someone to believe something. Though these are distinct, it is tempting to connect them. Dan, in saying 'Esau is hairy', explicitly expressed a certain thought, in doing which he may have w-expressed what Pia believes, wittingly or not. Which is to say, perhaps wittingly. If so, he may also have wished to call attention to the fact that he was doing so. One way of doing that, if means exist, would be explicitly to w-express Pia's belief, that is, to perform an act of doing so which, properly understood, was to be taken to be just that. Such would be an explicit w-expression of Pia's belief in the same sense that 'Esau is hairy' is an explicit expression of the thought that Esau is hairy. Perhaps words such as 'Pia believes that' just *are* a (conventional) means for doing so. In which case, had he said 'Pia believes that Esau is hairy', he would have done no more than w-express Pia's belief in a way explicitly to be taken as doing that. The grammatical object of his 'believes', on this view, would have referred to no object;<sup>7</sup> it would simply have expressed a thought. The truth of Dan's words depends only on their expressive part being a fair enough expression of Pia's belief, so on her being *somehow or other* such as to make them so. They mention nothing further to determine what else might express Pia's belief; or what else about Pia might make them fair expression. They contain no commitments on

<sup>6</sup> One need not assert things to express them in this wider sense. There is, e.g., the operator: 'Just to be clear, is what you're saying this: . . .?'

<sup>7</sup> Though its parts may have referred, as they otherwise would, to objects of any of indefinitely many sorts. A Fregean would still be free to say that it referred to a truth value.

these points. Nor do they suppose that there is just one specifiable way for Pia to be which would make them fair expression; nor do they pretend to identify any such way.

Frege's picture is different. If Dan said Pia to believe something, then, on it, he identified an object, B, such that he spoke truly only if Pia believably relates to B. B is identified by a semantics S. That something in Pia relates her believably to S, and so to B, is something else about her which makes Dan's words correct expression of her belief. Anything expresses what Dan said her to believe iff it bears some specified relation to S (such that, prefixed by 'Pia believes that', it would thereby refer to B). Dan's words commit him to, and their truth requires, all of that.

If Dan says 'Esau is hairy', he expresses a thought in speaking of nothing besides Esau and being hairy.<sup>8</sup> On the alternative to Frege, if he prefixes 'Pia believes that' to these words, he speaks of nothing new (except Pia). His commitments are thus far narrower than on Frege's view. This is not to deny that conclusions may be drawn from his words, if true. If Hugo now speaks words W, and if we grant that Hugo thus said just what Dan did (without the prefix), then we must also grant that Hugo too expressed what Pia believed. Whether he did, though, need not follow from the semantics of Dan's words. Nor need Dan have said what entails that such-and-such condition is the one Hugo's words must meet to do so. The result depends on what we would (or should) count as that which Dan did say, and on what we are prepared to count as saying the same as that. What Dan *said* leaves open, on the alternative, that there may be many sometimes correct ways of doing that.

The point of the above is only to indicate that Frege's picture is neither platitudinous nor required by anything obvious about what we do. Which is not to say that Frege is wrong.

### 3. Enumeration

What does Dan believe? Which beliefs are his beliefs? For these questions as they stand, we have not even the beginnings of an answer. For such a question in a proper setting, what would answer it is clear. Pia claims sales will rise in December, the holidays. Odile says they won't; the economy will nosedive. What does Dan believe? An answer might be: they will; they won't; you can't tell; it all depends on what you mean by rising; no opinion. Or, Dan has been said to think that Pia's yacht is longer than Odile's. There is the belief that, for some length L, presumed that of Odile's yacht, Pia's yacht is longer than L; and again the belief that, however long the yachts, Pia's is the longer of the two. Which of these beliefs is Dan's? Or does his contrast with these in some specific way?

<sup>8</sup> Or whatever his topics count, on an occasion, as being.

A background shows us our way about. Do backgrounds supply mere convenience, or are they rather something without which there would be no determinate condition on an answer? Do backgrounds select such a standard from a variety, each of which is sometimes correct? For Frege, the answer is clear. What a background affords is mere convenience. For any person and time, there is a definite set of objects to which he then believably relates—a particular selection from the objects there are so to relate to. These objects fix the right, though perhaps boring, answer to the question what Dan then believed; the same answer for any occasion for posing that question. To answer it, specify just these objects. A background merely excuses us from mentioning the boring ones.

If Frege is right, our ordinary answers to such questions are non-answers. Some say there is a tanker on the inlet; some deny it. What does Dan believe? That there is a tanker. Now we know, or so we think. But there are, on Frege's view, many objects of belief, all of which fix just those topics related just so. Did my answer fix one of these? Not at all. It selected only among the choices then on offer. Choices between different belief objects for the right topics were not on offer. So my 'answer' did not answer the question, even in part.

Suppose objects of belief served Fregean functions but were not fine-grained. Then I did (partly) answer the question asked. This does not seem right either. Sometimes, though not in the above case, we distinguish believing that there is, at that moment, a tanker on the inlet (glancing, perhaps, off the surface on its lunar trajectory) from the belief that there is, for a decent amount of present, a tanker on the inlet. Nor, outside a context, is there any place this game of distinguishing beliefs must stop. Any *given* coarse-grained specifying of a belief by its topics fails to make many such distinctions.

Conversely, if Frege is right, we normally regard many sensible questions as nonsense. Suppose that, on separate occasions, Dan, Greta, Hugo, and Pia all said that Esau is hairy. Perhaps all used the words 'Esau is hairy'. But each used them so as to have a different sense. Then if Pia believes, as we naively say, that Esau is hairy, there is the question: does she believe the sense of Dan's words, or the sense of Greta's words, or etc.? For Frege, such questions always make sense and have one right answer.<sup>9</sup> The question seems nonsense. What should one say for an answer? What distinguishing is wanted? We have said what Pia thinks of Esau; what more could one say?

Consider all that, at a time, *w*-expresses Pia's belief, and all that correctly ascribes it. For Frege, that which Pia actually does then believe—a unique set of items—makes this set of expressions *cum* ascriptions partition *uniquely* into equivalence classes. Each class consists of that which is certified (as genuine or correct) by exactly *one* thing Pia believes. Two expressions of belief express, *tout court*, the *same* belief of Pia's iff they belong to the same class. The classes of any alternate partitioning contain, *tout court*, items which do not express the same belief.

<sup>9</sup> Which is generally not 'Yes'.

## 4. Sentences as Models

Wittgenstein attacks Frege's picture in several ways. One emerges in the following:

The report is a language game with these words. It would produce confusion if we were to say: the words of the report, the sentence [*Satz*] produced, had a definite sense, and the reporting, the 'statement' grafted something on to that. Thus, as if the sentence [*Satz*], spoken by a record player, belonged to pure logic, as if there it had its pure logical sense, as if there we had the object before us that logicians treat. . . . The sentence [*Satz*], I want to say, has no sense outside of a language game. This is connected with the fact that it is not a kind of *name*. So that one could say: "I believe. . ."—it's *thus*', indicating (inwardly, as it were) what gives the sentence meaning. (*RPP* §488)

. . . The worst enemy of our understanding here is the idea, the picture, of a 'sense' of what we say, in our mind. (*RPP* §498)

The point is: no semantic item has a semantics which is *its* semantics no matter how it occurs or is applied. Any such item, even if it has some semantics as such, bears different semantics as applied or occurring differently. Only its occurrence-variable semantics fixes which standards for correctness, so truth, apply to it as so occurring. Fregean objects of belief violate this principle. Each has a fixed identifying semantics which, on any occasion, is just that semantics then determining what would express it, or ascribe it, and which ways for things to be would make it true.

The point can be put in terms of language games. Suppose that an item may be a move in a variety of language games—which is to suppose only that it may be spoken or expressed on a variety of occasions, for various purposes, within various activities. Each language game is governed by rules; its rules fix the standards of correctness (so, *inter alia*, truth) for any move made in playing it. Different games have different rules, hence different standards of correctness. Such an item, then, whatever its semantics as such, must be subjectable (compatibly with that) to different standards for correctness and truth, hence be capable of varying its semantics with the games it occurs in.

Language games are not to all tastes. Wittgenstein's point is also put by saying: sentences (of a language) are, in a key respect, the model for any semantic item. That respect is: sentences bear different semantics on different speakings, so occasions for them to count as bearing some. In particular, sentences vary their truth-involving properties. (That they do not bring the same set of these to every use is what Wittgenstein put as: a sentence in isolation lacks a sense.)

Any English sentence illustrates the point.<sup>10</sup> Consider 'There's a tanker on the inlet'. Suppose the inlet in question to be Burrard Inlet. Now suppose there

<sup>10</sup> Some will be attracted to other readings of the data. I have argued for this one in my 1981c, 1978, and Ch. 1 above.

is a tanker which, through novel propulsion techniques (hoverpower, say) is suspended about a metre off its surface. That sentence spoken of that situation, may well state truth—if, say, it is part of a report on local shipping traffic. But it may also speak falsehood. Suppose the inlet is abuzz with boats, all floating in the air above it. We are wondering when, or if, one of them will ever come down to sea. Dan speaks the sentence in this context. But he is mistaken. The only tanker still hovers. The sentence, so spoken, would be false. (Similarly, suppose I outfit my modest catamaran with a tank for transporting draught from the source to the inlet's far shore. If I am alone with it on the inlet and someone speaks the sentence, does he speak truth? On one way of viewing things, I own the world's smallest tanker. On another, it is no tanker at all, even if it has a tank. So my boat reveals a pair, at least, of things to be said of the inlet in that sentence.<sup>11</sup>)

Such contrasting pairs of cases show how the semantics of a given sentence as spoken on an occasion, including its conditions for truth so spoken, varies with its speakings. The same is shown by 'This wall is yellow'. Consider it spoken of a wall painted pale yellow. Does it matter to its truth so spoken that the paint covers white plaster? Again, there is no one right answer. For many speakings the answer is clearly no. Some speakings, though, may be properly understood as comments on the plaster—for example, where walls with rotting, yellowing plaster are being singled out for repair. For similar reasons, the sentence 'Snow is white' varies across speakings of it in what it demands for its truth so spoken. What it says on some speakings, but not what it says on others, is false if the famous motto 'Don't eat yellow snow' has point.

One could still suppose that a sentence's semantics remains constant across all its speakings. There is a price. One must then deny that sentences ever have such properties as being true, so seek other items, produced in sentences on occasion, which might be. With Wittgenstein, let us refrain from proliferating semantic items in this way. We may then say that a sentence sometimes *counts as* true, and sometimes does not. It so counts as spoken, or considered as spoken, on some occasions. In that case, its semantics varies across occasions for it to count as having one.

Wittgenstein's point is now that this feature of sentences is ineliminable; it cannot be refined away by shifting to other semantic items. *Any semantic item is intrinsically liable to vary the semantics it counts as having across occasions, if they may be multiple, for it to count as with some given such.*

Wittgenstein has an explanation for this principle which applies universally. Before stating it, it will be helpful to distinguish three sorts of semantic property. These need not be the only sorts. Central examples of the first sort are being true and being false. For further examples, choose some specific object—that chair, for example—or some specific situation—the one that obtains with respect to

<sup>11</sup> Or again, suppose the only item on the inlet is a war monument consisting of a rusting, overturned wreck of a tanker.

the inlet right now, say—or concrete state of affairs—the inlet's being as it now is, say. For each of these specific items, there are the properties of being true of it and of being false of it.<sup>12</sup> Call these *pure truth-involving properties*. Add to them properties of referring to, speaking of, or (to cover belief) being of such-and-such individual (but not of ways for individuals to be). Call the resulting category *pure referential properties*.

The second category is a superset of the first. It adds such properties as being true on condition that that wall be yellow, and being true of any situation which would now count as one of a tanker being on the inlet. Such properties need specifying in terms of truth (or falsehood). What distinguishes them from *pure truth-involving properties* is that they involve some sort of generalization as to which items would confer truth on what had the property: they speak of *any* concrete item that filled a certain general bill—that of being yellow, for example. They are not specified simply by some single item, or even a list of them, and whether it/they conferred truth on what had the property. Call these further properties *general truth-involving properties*. Call both pure and general truth-involving properties *truth-involving properties*. There may also be general referential properties. Speaking of whatever satisfies condition C would be one. Any property mentioned so far qualifies as a *referential property*.

Finally, there are what I will call *content-fixing properties*. These are specifiable without reference to truth, so are not truth-involving properties. But they are properties we generally take to be present and needed to fix just what a given item said or expressed or (where appropriate) otherwise represented as so. Such properties are, for example, saying something to be yellow, or speaking of being yellow, or ascribing being that to something, or saying there to be a tanker on the inlet, or predicating identity of an item A and an item B. I think that such properties are not equivalent to specific (conjunctions of) truth-involving properties. But the point need not be pursued here.

A semantics, on present usage, is merely a set of semantic properties. So imagine this. We collect together some of our favourite semantic properties. We have a semantics. We may then try to confer this semantics on something. We may, for example, invent an expression, E, and decree that E will have just that semantics. Our decree may or may not work. For our semantics may or may not be one that any item could have. Suppose, for example, that it included both the property of saying something to be a fish and that of being true of my computer. No expression could have both of those. Call a semantics no item could have *incoherent*, and use 'coherent' accordingly.

Incoherence may take subtler forms. On Wittgenstein's view, it may also be an occasion-sensitive affair. Consider speaking of being yellow, and being true of yonder wall. On certain occasions for describing the wall, the way it is counts

<sup>12</sup> Being true of an item is a tensed state of affairs; it involves being true of the item as it *now* is (while *that* property of being true of it is had).

as its being yellow. The above are a coherent semantics for what might then describe the wall. On other occasions, the wall does not count as yellow. Such is not what being yellow would then come to. No description of a wall could then have that semantics. A given semantics, then, might count as coherent, so available for an item, on one occasion; incoherent, so unavailable for any item produced on another.

This sometime coherence of a semantics explains why the semantics an item has, or counts as having, is an occasion-variable affair. Suppose that an item, W, counts, on a given occasion, as having some semantics, S. For another occasion, S may fail to be coherent, so a semantics any item could have. On such an other occasion W *cannot* count as having S. It *must* have whatever other semantics then fits it best.

An item has occasion-variable semantics if its semantics on an occasion includes pairs of properties which conflict or fit with each other in an occasion-sensitive way. Any item has such a semantics if it bears both content-fixing and pure truth-involving properties. So any item is of the right sort if it has content-fixing properties and is what might be true or false. Being a semantic item, it must then, on an occasion, have whatever truth-involving properties its content-fixing properties then count as requiring. But that varies with occasions. All our words, and all that we express in them, thus fall within this category. If an item has a semantics of this sort, then having its semantics occasion-independently cannot be secured for it by stipulation or like manoeuvres; it cannot be secured at all.

Take any content-fixing property and any truth-involving one it sometimes requires—say, being (now) true of A. These two properties cohere on an occasion only if the way A is then counts as being the way an item with that content-fixing property would say or represent A to be. If words speak of being yellow, then they count as being true of exactly what counted, on their speaking, as being yellow. No matter what the above two properties are, there is plausibly *some* occasion on which they would conflict. Occasion-dependence of semantics, though, requires less than this. All it needs is that for any content-fixing property, it sometimes conflicts with some truth-involving properties it sometimes requires. Which is to say only that the ways we may say things to be are such that different occasions impose different classifications of things as being that way or not. About this, there is little room for doubt.

How can a content-fixing property require or exclude a truth-involving one on an occasion? Suppose the property is saying something to be yellow. Then there is a rule in force: count an item with that property, on an occasion, as true of something just in case the thing then counts as being yellow. What would comply with that rule on an occasion? Wittgenstein is famous for speaking on this point. He reminds us, convincingly, that a rule regarded entirely in isolation requires or excludes nothing at all; *anything* might be compliance on some imaginable way of following the rule. A rule *may* require this or that on an occasion for

taking it to do so; it *will* do this if recognizing it to do so is part of our<sup>13</sup> natural reaction to the rule in those circumstances. The occasion-dependence of semantic coherence then follows from this: our natural reactions are liable to be sensitive to the circumstances in which the rule is to be reacted to. Our actual treatment of words demonstrates this sensitivity: the manifest occasion-variability of words' semantics is an effect, so proof, of Wittgenstein's point about the limits to what a rule could, in principle, require.

Wittgenstein establishes his point in *PI* §§84–7. He raises the topic again in *PI* §§185–202, not to establish a new result about rules, but rather to apply it, specifically to remind us of issues of semantic coherence. Such is important for the topics he there discusses: meaning and understanding. It is similarly important for belief. If for Pia to believe that Esau is hairy is for her to be in some fixed way for a duration, it matters that only so much *could* be fixed about her, and about what she thus believes, apart from our reactions in particular circumstances to the semantic consequences of being the way she is. In semantic matters, our reactions are notoriously occasion-dependent.

## 5. Media and Belief

In one respect, sentences may serve as a model for any semantic item. In others they do not. A Fregean object of belief would differ from sentences in crucial respects. Frege would deny the similarity stated in the last section. He equally overlooks those differences.

*Sentences.* A sentence is a medium, or vehicle, for saying or expressing things. Since it aims only to be *usable* for saying things, there need be nothing in particular that it itself says. Though 'says . . .' sometimes applies to sentences, for limited fillings of the blank, if there were anything the sentence itself said to be *so*, that would spoil its role as a medium.

Because a sentence is but a medium, we can afford considerable insouciance as to its truth-involving properties. If a sentence is not true, or untrue, that no more displays a failing on its part than that lack would exhibit failure in my raincoat. We do not take a sentence, on its own, to aim for truth; nor does it, since there is nothing it says on its own to be so. Nor is it a flaw if a sentence fails to have a fixed set of truth-involving properties, such as imposing some definite condition for truth. Such a condition would be called for only if the sentence as such aimed for truth, as it does not.

A sentence often acquires specific truth-involving properties on a speaking of it. We can afford to let these be whatever its constant semantics and the occasion

<sup>13</sup> I purposely leave this vague. What is in question, roughly, is what we are prepared to recognize as the reaction of a reasonable judge. For further discussion see my 1989, ch. 2.



of that speaking jointly require or make reasonable. There is no reason why these properties should not vary considerably across speakings. To encapsulate this in a slogan: for a sentence, truth-involving properties follow wherever other semantic properties may lead.

We care quite differently about content-fixing properties of sentences. We expect some such to attach to a sentence independently of its speakings, and to come with it into every speaking (spoken as meaning what it does). It is these and the occasion, we hope, that fix sufficient variable truth-involving properties. The words 'is white' speak of being white. In terms of this we hope to see what 'The wall is white', on a speaking, requires for its truth. (The sentinel, Bing, reports 'The Wall is white', referring to the Great Wall in the distance, as it reflects the morning sun. It would not matter for his words if the wall looked muddy brown when examined close up. What would you expect to count, in Bing's situation, as being white? It would matter a lot if the slate-coloured sky had communicated its dull grey to the wall.)

What is fixed in a sentence's semantics allows us to take it correctly to mean such-and-such, and to say what it and its parts mean. Meaning such-and-such is a fixed (speaking-independent) feature of a sentence, and exhibits what is so fixed in its semantics. It is possible to overstate this case. Radical enough changes in circumstances may change what E could correctly be said to mean. At any given time, though, there is something it does, and may truly be said to mean; which fixes a feature it brings to any occasion for then speaking it.<sup>14</sup>

There is a moral here. An expression's meaning such-and-such is a feature it is meant to carry into every application. There being such a thing as what it does mean then requires sufficient constancy (occasion-independence) in enough of the right semantic properties. The right ones are content-fixing ones. Such constancy generally obtains. But it is not independent of our insouciance about an expression's truth-involving properties. Truth-involving properties, particularly pure ones, are intrinsically liable to come into occasion-dependent conflict with content-fixing ones, as above illustrations show. For sentences, there is a clear resolution to such conflicts: content-fixing properties lead; truth-involving ones follow. Such allows content-fixing properties to remain constant across conflicts. They do so *only* due to our insouciance about those others. Our ability to state meanings thus depends crucially on this insouciance. Expressions may be supposed to mean this or that precisely because they are *not* supposed to have any fixed set of truth-involving properties, notably not pure ones. Stability in their content is *at odds* with stable truth-involving semantics.

Meaning is useful. It provides a shortcut around direct appeal to the specifics of an expression's history in its language. If we know what the words 'The wall is white' mean, then the circumstances of their use on an occasion provide a basis on which to judge their semantics as so spoken. What we need to know

<sup>14</sup> In properly speaking its language.

about the use, roughly, is what was being done, or what was properly taken to have been being done, in so speaking them. Imagine for a moment that we could not refer to what the words mean in establishing their semantics so spoken. Presumably, what they do mean depends, in some complex and doubtless not precisely specifiable way, on the ways in which they already have been used in their language: on the expectations they have aroused in prior hearers of them; what their prior users meant them to accomplish; how they have been treated, by recognized authorities and others, in assessing their correct and incorrect uses; and so on. We might then appeal directly to these facts. Without appeal to meaning, we could do little else. With grasp of meaning, we can usually omit such excursions into history.

*Belief.* Belief is not thought. Believing that Esau is hairy is utterly different from thinking 'Esau is hairy'. (We typically believe what we think, as we typically believe what we say. The one fact ought to mislead us into confusing thoughts with belief no more nor less than the other makes us confuse words with belief.) Nor is belief a medium for anything. Not surprisingly, we take a very different interest in belief than we do in sentences. To begin with, what someone believes depends crucially on when, or of what, it would be true. By contrast, neither what a sentence is (which one) nor what it means depends on this. If Pia believes that Esau is hairy, then it is surely essential to what she thus believes that certain specific smooth states of Esau would make her belief false. If Bing, on that sunny morning, believed that The Wall was white, then it is no good saying him to believe just that, never mind when it would be true. It matters to what Bing believed that the muddy brown of the stones of The Wall viewed close up would not tend to show him wrong.

A belief may, exceptionally, turn out to be neither true nor false. (Some would say it is then no belief. I am not among them.) Dan believes, as one might put it, that there is a tanker on the inlet. The only one hovers a metre above it. Hoverpower is a surprise. Given what Dan took to be so, and the circumstances of his doing so, the actual state of things may not justify labelling his belief true; nor false. (Similarly, perhaps, for Pia if Esau turns out to be depilated.) If this does happen, though, things are not as they should be. There is a failure of some sort. The belief has not failed of its purpose, or fallen short of its aim, because beliefs *have* no purpose or aim. The world has failed to co-operate as that belief (pre)supposed it would. A belief's being true or false presupposes that the world classifies determinately on a particular way of doing so. Such suppositions are substantive, so sometimes fail to hold. Beliefs presuppose such things, as sentences do not.

Similarly, if what someone believes is identified (in part) by some set of (pure) truth-involving properties, that set need not, and I think never will, be maximal. There may always be further properties that might have been joined to the set—properties of being true (false) of such-and-such—without rendering the

result inconsistent. There may always be truth-involving properties, compatible with all those which fix what is believed, which determinately neither fit nor fail to fit that belief *tout court*, but which might count as doing either on some occasion. If Esau were hairy for three minutes out of every ten, smooth for the rest, what must we say of Pia's belief then? Perhaps not either that it is true or that it is false *tout court*; but perhaps, on some occasion, the one thing, and on others the other.

While what identifies what is believed need not be a maximal set of truth-involving properties, it should be a sufficiently rich one: rich enough to satisfy us that we know what is believed. Where it is not, we cannot take there to be something in particular the believer believes. This notion of rich enough is, of course, occasion-sensitive.

Let us now suppose, with Frege, that one believes an object.<sup>15</sup> Then that object must be rich in truth-involving properties, including pure ones. For it can take no less to identify it than it does to identify what is believed. Those properties are part of what individuates the item. So, as with what is believed, a rich set of these adhere to it permanently—not just in the temporal sense, but in the quite different sense that the item has them occasion-independently.

Despite the above, we typically say what someone believes in terms of content-fixing properties, and not in terms of truth-involving ones. We say, for example, that Bing believes that The Wall is white. We thus identify his belief in terms of The Wall and being white—what he says the wall to be—and thus in terms of the content-fixing properties of speaking of and ascribing those things.

This practice serves an intuition. It is that beliefs, like words, cannot *simply* be true or false (of this or that); but if they are so, must be so in virtue of something else about them. Content-fixing properties are *just* the right sort of something else. If Pia says 'Esau is hairy', then her words are false given Esau's smooth condition. That is because of what her words said him to be: hairy, where being that is to be understood as the antithesis of being smooth. *Mutatis mutandis* for believing that Esau is hairy.

On Frege's view, there is an item which is the object of Pia's belief that Esau is hairy. Its semantics must identify what it is that Pia thus believes. Respect for the above intuition then requires that the semantics of that item include appropriate content-fixing properties. (It would not matter even if there were truth-involving properties which fixed just what the right content-fixing properties would be. In that case, a belief object, in having the former, has the latter too.) To sum up, a Fregean object of belief must be rich enough in truth-involving properties, including pure ones; but it must also have sufficient content-fixing ones.

Wittgenstein's insight about semantics shows a consequence of the above. To have the required sort of semantics, a Fregean object of belief must have just that mix of properties—pure truth-involving ones and content-fixing ones—which

<sup>15</sup> Or believably relates to one such.

conspicuously conflict in occasion-sensitive ways. So, if a Fregean object of belief has the semantics it needs to, it must bear its semantics occasion-sensitively. It will count as having different sets of semantic properties on different occasions for it to count as having a semantics at all, notably ones for expressing it.

Any item with content-fixing properties always counts as having all the truth-involving ones those content-fixing ones require, and as lacking all those those content-fixing ones exclude. An object of belief, if an object, purchases determinacy in what is believed by bearing stably a rich set of pure truth-involving properties. If it does that, then its content-fixing properties cannot be invariant. Such an object thus reverses the situation for sentences: for an object of belief, truth-involving properties lead; content-fixing ones follow where they must.

An English sentence has, as such, a semantics. Neglecting ambiguity, it carries this to, and on, every speaking, and supplements it on each. The sentence, say, speaks of being white. Then it does that on every speaking, adding there just the truth-involving properties that that feature and the occasion require. A Fregean object of belief could not have such an occasion-independent semantics. It could not have, as its 'semantics in isolation', a semantics it carries on every occasion for expressing it, or saying what it is. Moreover, its content-fixing properties cannot be so invariant. There is, then, nothing for an object of belief parallel to meaning for a sentence; no occasion-independent specification of its content. Such means, for a sentence, provide a shortcut around appeal to history in recognizing what an occasion requires of and for it. For a belief object, there is no parallel shortcut around appeal to the believer's life in seeing what an occasion demands for expressing what he believes, or, with Frege, that object he believably relates to.

We cannot have what Frege supposes we can. Here what he supposes must be carefully distinguished from what we surely can do. On occasion, we can specify what Bing believes; sometimes in saying 'Bing believes that The Wall is white'. That is the typical and preferred way of doing so. On it, we specify belief in terms of content-fixing properties—in this case, speaking of being white and saying The Wall to be that. The above does not deny this. But in doing such things we pretend to no more than identifying belief occasion-dependently; saying what counts, on the occasion, as what someone believes. We do not suppose the features we produce to count on every occasion as features of his belief. Bing's camp may harbour sceptics about the way the wall looks. One might ask what Bing believes. The above would be a fair and true reply. Suppose we are among a group of people examining the wall closely. Being bad at colours, they are disputing whether it is really beige, or rather grey. One might ask what Bing believes. We *could* reply, 'He believes that The Wall is white'. We would not thus speak truth. That is not what he here counts as believing about The Wall. This in no way reflects on our earlier report. It denies nothing said in giving it.

Our belief cannot be expressed in the same terms on every occasion for doing so. That is a point about *expression*. It converts to one about *what* we believe.

Whether what Bing counts as believing about The Wall is that it is white varies with occasions for Bing to count as believing this or that, and what, on them, The Wall's being white would come to. If what we believe are objects, then either on different occasions we count as believingly related to different ones, or any such object counts as bearing different semantics on different occasions for expressing it. In neither case does the object as such exhibit what is fixed about someone for a duration in his believing such-and-such.

The above point has like-sounding, but distinct, relatives. Consider a suggestion of Gareth Evans. Evans remarks:

the changing circumstances force us to change in order to keep hold of a constant reference and a constant thought—we must run to keep still. (1985, p. 308)

The remark is prompted by examples like this. On a given day, someone might w-express a belief in saying, 'Today is fine'. The next day, those means for expressing that belief are unavailable. Stronger, *no* words could then have exactly the semantics of those words so spoken. If the same belief can still be w-expressed, as Evans suggests it can, then necessarily an expression of it will have different semantics from any expression of it the day before. Evans's point is about *expressions* of belief: what would express a given belief varies with time (and perhaps with place).

I do not think Evans sees his point as transferring to the semantics of beliefs themselves. For someone who believed, and continues to believe, what yesterday was w-expressed in saying 'Today is fine', Evans seems to suggest an occasion-invariant which identifies, *tout court*, both what is believed and a constant semantics for it (though it is unclear how such an invariant might be specified). But these are issues for another time.

As for expressions, the present point may be a natural generalization of Evans's. But it *is* a generalization, and not one Evans would accept. Occasion-variability is not the same thing as variability across time or place. Evans speaks only of the latter. Fix a time and place, and the variation he detects in what would express a given belief is at an end. The belief's constant features plus specific values of those variables fix uniquely what expresses what.

The present point is not that. An indefinite variety of occasions or speakings might occur at any given time and place. On each a different semantics might be required to fix *what* Bing believed in relating to his environment as he did at that time. In camp, some doubted his report, an occasion for expressing his belief in given terms. At the same time and place, there could have been people debating the look of the wall close up. With no change in *Bing's* situation, different terms would then be needed for expressing Bing's belief. That occasion-variability exists suggests no variables, values of which would fix uniquely what expresses what, or has what semantics; nor that there are any such variables to discover.

Frege preserved the healthy intuition that beliefs are what may be true or false. A recent proposal suggests otherwise. Its aim is to capture the implausible idea

that what we do believe is what we would have believed no matter how radically our environment might have been otherwise. Pia believes that she is drinking tea. On Frege's view and mine, what she believes is (partly) fixed by its topics: drinking *tea* is what she believes she is doing. These topics, in turn, determine whether her belief is true. Had there never been such a thing as tea, Pia could not have believed what she does. Still, some insist, she would have related believably to the same object she now does. The proposal is then to regard this object as a function taking as values what Frege or I might recognize as what someone believes. The arguments for such a function would be environments, or possible ones: the situation the believer happens to be in, or specific aspects of it. The present point adds to the demerits of that idea. Fix the environment Pia is in, and the function her belief is meant to be should take on just one value. But given the present point, there is no such value for such a function to take on. Without altering *Pia's* circumstances, her belief counts as having different semantics on different occasions for expressing or thinking of it, even at a given time. On different such occasions, different things count as what it is she (relevantly) believes. That is the clearest proof that there is no such function.

Fix all else about a person at a time; his conditions both inner and outer. His belief, inclusive of *what* he believes, wears a semantically different face, takes on semantically different forms, as viewed from different vantage points, or occasions for beholding it. That is shown by the occasion-dependence a belief's semantics would need to have. There is no particular semantics for the conditions of a believer to map onto. That is why a belief cannot be represented by a function. The shifting semantic form of a belief suggests that which words ascribe belief correctly, and what is required for any given such to do so, has at least as much to do with the position of the ascriber as with the condition of the ascription's subject. Which suggests such words need say no more than relates their subject to what is said in speaking them, nor identify more fixity in the subject than that.

Frege's picture is mistaken. Believing such-and-such is, to be sure, being in some fixed way for a duration. Frege represents that way as relating believably to a given object; where that object identifies that fixity in the way one is. The object is individuated by a semantics, which it thus bears occasion-independently. This semantics identifies *tout court* one thing to be believed; that it thus identifies what he believes is a fact, *tout court*, of the believer for that duration. Frege is wrong because no semantics capable of identifying something to be believed or not could be the semantics of anything occasion-independently. So having that semantics could not be a fixed feature of what is believed; nor believing what has that semantics a fixed feature of the believer for the duration. The first two points in Frege's picture have now run into trouble. The last point has played a role: since on it we refer to objects of belief by their semantics, such semantics must be specifiable. The controversial third point has played no role.

## 6. Remnants

On occasion we say what someone believes. Or express it. If, on one occasion, we are told what someone believes, we can often tell, on another, how, and how not, to express that. There is less in these facts, or in anything else we clearly do, than Frege's picture contains. All hope for that picture may not yet be cut off. But if it can be shown not to be required, the ground is then clear for a better.

If Frege's picture seemed what grammar requires, it might seem worth accommodating to a world of occasion-dependent semantics. There are two tacks to take. One is to preserve the idea that for each thing one believes there is a particular object such that to believe that is to relate believably to just that object for the duration, and that this object represents what is fixed about one in believing that. Such an object must then bear different semantics on different occasions for expressing it. The other tack is to allow that, for a given thing we believe, we count, on different occasions, as relating believably to different semantic objects in virtue of that, though just one object for each occasion. Such an object *might* be individuated by its semantics, since it is called on to function on but one occasion. This amounts to the thought that for each thing we believe there is, on each occasion, a (unique) semantics which performs for it the certifying functions a Fregean object of belief was meant to; together, perhaps with the idea that there is a unique correct way of counting things we believe. I will suggest that the first tack, if distinct from the second, is incoherent; the second one adds pointlessly to what we clearly do.

Consider the first tack. Since we can say what people believe, it must be possible to say which object one relates to in believing thus and so, and, for any such object, which one it is. Since, on the picture, our belief ascriptions refer to them, such reference must be possible. How might we do it? There is an obvious suggestion. On a given occasion, we said what Bing believed in saying 'Bing believes that The Wall is white'. If we spoke truth, then that fact identifies something fixed about Bing for the duration: he was such as to make our words, on their occasion, a true expression, and ascription, of his belief. On the first tack, our words 'The Wall is white' *refer* to an object of belief. Which object? The one that certified those words as an expression of it on that occasion. What I just said, then, specifies one such object.

But *does* that device refer? Could there not be several such objects, each of which would certify that same expression of it on *that* occasion? If so, nothing so far distinguishes one of these as referred to above. So we may not suppose that the above is meant for referring to such things.

A Fregean object of belief is, after the accommodation, an item that bears different semantics on different occasions. So it must be possible to say which *single* thing it is that, considered on different occasions, counts as having different

such properties. What might identify such an item as the same one considered twice? What sorts of properties might it have other than the ones that may vary? What must be held constant for *it* to vary its semantics?

To treat such objects as varying semantically across occasions is to model them on sentences, which bear different semantics on different speakings. But there are clear ways of identifying sentences independently of such variable aspects of them. It is clear what sorts of *non*-semantic properties they might have. Spellings, for example. Beliefs are not media; nor vehicles. Which limits the non-semantic properties they might have. In terms of what constants might such an object be specified?

There are now two problems. One is that we do not know what constants might identify an object of belief. Nor have we any conception what it might be to discover them. The second problem is more serious. Grant that such objects are individuated by such-and-such features. There is now no guarantee that there could not be two such items, identified by different such features, which coincided on a given occasion in certifying the same things, and, perhaps, just the same things, as expressions of them. As long as our access to such items is only by their effects on an occasion—that is, by what then expresses them—there could be no such guarantee. So if we do conceive of such items as identifiable other than by their semantics on an occasion, we may not yet suppose that we have means of referring to them.

That we can view such items only *via* their expressions is part of the content of the thought that these items provide no shortcut around the biography of the believer, of the sort that meanings provide for the histories of what have them. That Pia relates believingly to such-and-such objects is only to be grasped, on an occasion, by seeing what anyway then would express her belief. That is not shown by what given semantic objects require; it is shown only by a direct view of what her life commits her to.

Beliefs not being media, it is tempting to distinguish them no more than need be from their semantics on an occasion. For example, they might be thought of as functions; *roughly*, from occasions to semantics. For a given occasion as argument, the value the function would take on would be just that semantics which certifies as such just those expressions which then express that belief the function is to be supposed to be.

The idea prompts three remarks. First, that there are such functions is no platitude. There is a determinate function only where it is determinate what its arguments and values would be, and, for each argument, that there is just one value for it to take. There are such functions as just described only if, for each thing one might believe and each occasion, only one semantics could count as certifying (in a given way) just those things that would then express that; so only if there is just one way of counting things, on an occasion, as expressing that or not; and only if specifiable features of occasions are what always decide the right semantics for them. Occasion-variability in what would express a belief was



demonstrated by expressions which, on one occasion but not another, would express a given thing someone believed. The conclusion was that the same belief must sometimes be expressed in different terms. It does not follow that there is any set of specific variables in circumstances values of which regularly fix what would express what (or even what would express thus and so). There is no reason to think there are. Conversely, for any given proposal, it is no platitude that there are functions from those variables to (semantics for) the proper expressions, wherever given values of them hold, of a given belief.

Second, Fregean objects of belief must be specifiable; so such functions must be. They may not be thought of merely as indefinitely large and variegated sets of ordered pairs, which are somehow just 'there' even if we cannot say which sets they are. Which means that, for each function, there must be a way of saying what it would do for any argument. Frege's notion of function is really what is wanted here: a function is a specifiable method of determining values for any argument (see Frege 1891, esp. pp. 26–7 in trans.). (Even if we then add that two different methods cannot *always* determine the same things.)

Third, the same problem now arises again. However exactly such functions might be specified, could there not be several each of which takes on the same value for *some* argument? It is hard to see why not. But then we still lack means for referring to Fregean objects of belief.

The original thought was: suppose that 'The Wall is white' would, on some occasion O, w-express Bing's belief. Then we may speak of 'the object that certifies that expression as such' (via some given relation between its semantics and the expression's). Stronger, there is a uniquely determined set of belief expressions which are just those that, on O, would w-express that belief the above expression did, thereby expressing Bing's belief. We may speak of 'the object which certifies just those as such'. So there is just one object which certifies the given expression, and it also certifies just the right set of others (if there are any). Though we have no right to stipulate such things, let us now suppose them anyway. In that case, what any such object does on an occasion depends on nothing about it other than the semantics it counts, on that occasion, as bearing. (Since there could not be two such objects that did that.) Moreover, what the object does on any occasion is fixed by its semantics on any other: since nothing else matters about it on an occasion, its semantics on the one must be predictable from its semantics on the other. The first tack is now transformed into the second, assuming no more than that does. So let us turn to this second tack.

One thing the second tack supposes is that, for any person, time, and occasion, there are generalizations to be made, in semantic terms, as to what would then count as expressing that person's belief at that time. Some semantics, for example, might fix an equivalence class of enlargements of it, such that any member of

the class is a semantics for what would then express such belief. Similarly, there are generalizations to be made as to what would count, on an occasion, as w-expressing some one thing to believe. So far, nothing is controversial.

To this non-controversial core, the second tack adds some or all of the following. First, any two expressions of belief on an occasion either do or do not, *tout court*, count as then expressing the same belief. In such matters, there is just one right way of counting. That way is captured by just one set of generalizations about what would, on an occasion, express someone's belief at a given time; one generalization for each thing to be believed that that person does believe. Second, Frege's picture supposes that when we ascribe belief to someone, we actually *make* such a generalization. For on it, in doing so we refer to a particular semantic object, thus to a semantics which captures exactly one such generalization. We thus commit ourselves to exactly those things the generalization certifies as also expressing that person's belief, and the same belief as that we ascribed. So, going beyond our own (embedded) w-expression of a belief, we take a general stand on what would, and what would not, express the right thing. Third, it may be assumed that a semantics which captures a correct generalization for one occasion also determines, for any further occasion, which semantics would then capture the right generalization as to what, on that further occasion, would express the same thing-to-be-believed as what the original semantics identified for its occasion.

These extra assumptions are not required by anything we clearly do. That is seen by observing that the alternative proposed in section 2 does not make them. On that view, in saying 'Esau is hairy', we w-express a belief. In saying 'Pia believes that Esau is hairy', we just explicitly represent ourselves as w-expressing Pia's belief. Our expressive words get no new topics; they all still speak only of what they otherwise would have. We thus express no further commitments as to what would express her belief, either on the occasion of our speaking, or on any other. We identify no way of counting things as saying or expressing the same, though if our ascription is true, then whatever does count as expressing the same as we did also counts as expressing Pia's belief. Our ascription thus leaves it open that there are many sometimes correct ways of counting such things, depending, perhaps, on our interest in doing so, even given that our interest is directed towards expressions on some given occasion.

It is also left open that whether given other words express Pia's belief depends not only on judgements as to whether anything with this semantics expresses the same as anything with that one, but on specifics as to what being true to Pia's life requires. We sometimes cannot state Bing's belief by saying 'He believes that The Wall is white'. That identifies terms we may not use; but does not yet identify which terms we *may*. May we say him to believe that it glistens white in the sunlight? Or that it counts as white on such-and-such way of counting

things as being white or not? Any such formula may fairly represent what Bing relevantly believes about The Wall. But it may also fail to.<sup>16</sup> Whether it does depends, plausibly, on further particularities of Bing; not merely on the fact that our original expression, with its semantics, was, on its occasion, a fair statement of his belief.

Frege saw no problem in referring to, so identifying, objects of belief: our referring words have a (normal) semantics; that is the semantics of the item we refer to; having it identifies the item as that one; it does so occasion-independently. Frege missed a metaphysical insight: for any way for things to be, so any way they can be said to be, there is, in sublunary affairs at least, more than one sometimes correct way of counting things as being that way or not. The semantic consequences of that metaphysical fact scotch Frege's idea. His picture of belief is left either without sense or without point.

Grammar may seem to require Frege's picture. 'Believe' being a state verb, belief has duration; so to believe thus and so is to be in a fixed way for a duration. Belief has content: one's belief varies with what one believes. So what one believes seems part of what is fixed about one for the duration. Further, belief has semantics: one may believe *truly or falsely* that such-and-such. Frege's picture is true to all this.

So, though, is the alternative. On it belief ascription aims, not to *identify* an occasion-independent aspect of someone's believing thus and so, but to *express* what he believes in occasion-dependently proper form. I express Bing's belief on occasion in saying 'The Wall is white' in saying 'Bing believes that The Wall is white'. There is then something fixed about Bing for a duration: his being such as to make my words justly express his belief. There is no pretext of saying more as to what about him makes this so; or that any *one* thing is always what would. Bing thereby relates to *some* semantic item: my expressive words. That relation is: they said what he believes. Their saying what they did no doubt depends on their having the semantics they did. But what Bing believes, so believably relates to, is what they said; *not* their semantics (since it is not their semantics they speak of). What they said, and his relation to it, accounts for the semantics of his belief: if they said such-and-such, then what he thus believes is true iff it is true that such-and-such. (Here talk of truth requires no mention of a semantic item.) My words did not *say* what it was they said. Which leaves it open that many substantively different ways of identifying what they did say, and what else would say that, might each sometimes count as correct.

Frege's picture suggests that what Bing believes depends on precisely how he represents relevant topics to himself: his conception of The Wall, or what he

<sup>16</sup> It may, e.g., be unfair in particular circumstances to represent Bing as with explicit views on how things might be counted as being white or not.

is prepared to count as being that, and so on. On the contrary, what words say is typically identified by their topics: the objects and ways for things to be they spoke of. If so, then my expressive words purport to identify just one fact about the way Bing represents the world. As that fact sometimes is put correctly: that he represents The Wall as being white. Belief ascription is thus confined to reasonable ambition. Working oneself into a position where one's expressive words would have just that semantics that captured *all* of the way Bing represents to himself relevant items and ways for them to be is nearly always a blatantly impossible undertaking. Actually saying what someone believes ought not require aping Pierre Menard aping Cervantes. Nor is there call for special linguistic resources, like 'believe', for that express purpose.

There may be many substantively different ways of saying what words' topics were, so what they said, so what would say the same. Different occasions demand different ways of counting things. Where words spoke of an object, though, as opposed to a way to be, saying what they did demands speaking of that same object. The sort of object it is constrains the possible ways of counting a thing as being that object or not; correspondingly the possible ways of counting words as saying what those did. On the alternative to Frege, words in belief ascriptions always maintain their 'normal' reference and only that. Words that do that may speak of any of many sorts of objects; anything, in fact. On Frege's view, words prefixed by 'N believes that' always speak, as a whole, of *one* object of one given sort; moreover, an object subject to just one possible way of counting things as being it. So there can be just one way of counting other words as saying what such prefixed words did. Which constrains our ways of speaking of belief *much* more than practice suggests them to be constrained.

In assigning a semantic object to belief as its object, Frege represents someone's believing that A as a semantic state: a state specifiable by that semantics which identifies, *tout court*, what it is belief in. Though not itself a mental state, such a semantic state might repay empirical investigation. For, as Frege notes (1918–19, trans. p. 26):

Although the thought does not belong with the contents of the thinker's consciousness, there must be something in his consciousness that is aimed at the thought.

That something must distinguish the semantic object thus targeted from every other. So it must be a sort of sense-surrogate: for each (semantic) feature by which the targeted sense might differ from any other, there must be some other feature of the thinker's mind (or brain) that identifies the 'aimed-at' sense as one with that feature, rather than some other. To reject Frege's picture is then to identify a sense in which belief is *not* a mental state: it is not a state semantically identifiable as Frege thought. So there is no semantics for which there is a

surrogate to be discovered. Here is one point at which, as Wittgenstein reminds us, what empirical discoveries there are to make is *not* a question grammar leaves open.<sup>17</sup>

Postscript: *Was* Frege wrong? As I now see it, the answer turns on whether he really thought there was a unique right way of counting *Sinne*, so thoughts. I am now unsure of this.

<sup>17</sup> My present conception of the phenomena discussed here and their point has been influenced by (sometimes critical) comments by John Campbell and David Hunter.

## On Concepts of Objects

How does reference relate to truth value? How, that is, do the referents of proper parts of a whole with truth value—the objects, properties, relations, etc., those parts speak of—relate to its truth value? Frege held that the referents of proper parts fix the truth value of their whole uniquely. He recognized apparent exceptions—so-called opaque contexts such as indirect quotation. But he made these no exception by finding different referents for words in such contexts. Bracketing such cases, his view means this. Insofar as words have senses, or bear particular understandings, once their referents (what they speak of) are fixed, such senses, or understandings, can have no further consequences for their truth value. Vary sense with reference fixed, and truth value remains constant. This makes it harder to demonstrate that words *have* senses. For there can then be no pairs which agree in reference but differ in truth value—the clearest way in which sense could make its presence known.

Frege's view recommends itself in either of several ways, depending on how we approach reference. On one approach, it is simply healthy realism. Take a simple case. Someone says 'The duck is dead'. Moving from those words, or that thought, to referents, we arrive at a certain duck and a certain way for ducks to be: dead. As long as we arrive at those referents, we arrive either at a duck that is that way, or one that isn't. (Forget for the moment brain-dead ducks on life-support systems.) The senses, or understandings, if any, which we meet along the way, and which, perhaps, help us along the way, cannot change what we thus arrive at. If the duck is that way, then the words, and the thought they express, are true. If the duck isn't that way, then they are not true. Given the healthy thought that senses cannot change the way the duck is, there is thus no room within these parameters for sense, or understanding, to change truth value. The point stays as good as it is as the complexity of words increases.

On a different approach, Frege's view seems not in need of argument—in fact, not exactly a view—but true by legislation. We may often say what it is to be a referent in terms of effects on truth. The referent of 'Pia' in a given 'Pia is asleep' is just that individual (if any) such that the truth of those words turns precisely on whether that individual is asleep. Generalizing, the referent of a proper part of a whole with truth value would just be whatever it is on which the truth of the whole depends, given the contributions of the other parts; whatever it is that determines that part's contribution to determining when the whole would

be true. So to have the referents of all proper parts is to have everything that determines when the whole would be true.

On the other hand, this view of Frege's is one of the later Wittgenstein's main targets, notably in the first 120 paragraphs of *Philosophical Investigations*. The point of the discussion of language games, with which the work begins, is that naming, or reference, even in the present broad sense, underdetermines conditions for the correctness of wholes, notably, where relevant, conditions for their truth. Wholes with given referents, embedded in different language games, would be true under any of many very different sets of conditions. What Frege leaves out on Wittgenstein's view—in fact, explicitly denies—is that 'true' and 'false' are in the first instance evaluations of particular historical events—speaking of words on particular occasions, in particular circumstances—and of the fittingness of the words for those circumstances. Those issues do not reduce to facts about referents; nor can we make them do so by definition. There is nothing to play the role that reference would thus be assigned. If Wittgenstein is right, then the healthy realist thought that senses cannot change the condition of a duck must be made to live in peace with the equally healthy thought that the truth of our thoughts about ducks, or whatever, depends, *inter alia*, on what we make of them.

I think Wittgenstein was right. My main reason concerns what is left open by references to ways for things to be—properties and relations. We may take an English predicate—'is green', say—identify what it speaks of—being green—and still identify many distinct things to be said in using it to speak of precisely that. (I will say more about this in a moment.) But it is also so, as I will argue, that a term's having referred to such-and-such object is compatible with its having made any of many distinct contributions to the truth conditions of its whole, so that two wholes with parts thus alike in reference may nevertheless differ in their truth value. Such variation manifests itself, e.g., in identity statements—statements whose truth depends on their referring to one thing twice. In this case it is difficult to make the Wittgensteinian view even coherent. I hope to do at least that much here, and to suggest that, once coherent, the view is eminently plausible.

## I. Identity

Consider a pair of identity statements—say 'A is B' and 'C is D'. Suppose they agree piecewise in reference: both speak of identity and for some U, both 'A' and 'C' count as speaking of U, and for some V, both 'B' and 'D' count as speaking of V. On Frege's view, the two cannot differ in truth value. If we dismiss the possibility that both may lack a truth value, we must then say what Frege did say:

If the sign *a* is to designate an object for us, we must have a criterion which decides in all cases whether *b* is the same as *a*, even if it is not always in our power to apply this criterion. (Frege 1882*a*, trans. p. 73)

There are two distinguishable ideas here. The first is that for every object, *a*, there must be a criterion for being it—a ‘criterion of identity’. Such a criterion speaks to certain questions: for any object, *b*, ‘Is *b a*?’, for any property *P*, ‘Could an object with *P* be *a*?’ It says something about what it would be, or what is required, for an object to be *a*. Right or wrong, that idea is not in question here. The second idea is that every object needs a criterion which *answers* every such question, or, more modestly, is consistent with at most one answer to it, ruling out all others as wrong *tout court*. I will call such a criterion *categorical*.

Does an object need a categorical criterion? Must there be unique answers to all the questions such a criterion would address? There are reasons for thinking so. First, there is healthy realism applied. Beginning with a statement of the identity of *a* with *b*, we move from words to referents. Either our terminus is one object, or our termini are two. We cannot move to *objects* which are either one or two depending on how you understand that statement. That idea makes no sense. So, it seems, all statements of the identity of *a* and *b* must share a truth value. Which means there is a unique right answer to the question whether *a* is *b*. Second, there seems to be an argument that this must be so. For suppose two such statements differed in truth value. They might, it seems, be expressed thus:

1. A is<sub>1</sub> B.
2. C is<sub>2</sub> D.

Here ‘A’ and ‘C’ refer to one thing (*a*), and ‘B’ and ‘D’ refer to one thing (*b*). The subscripted ‘is’s may differ as much as possible, *provided* that (1) and (2) genuinely state identity. Suppose that (1) is true and (2) false. Then (3), the denial of (2), is true:

3. C is<sub>2</sub> not D.

Since ‘A’ and ‘C’ refer to one thing, Leibniz’s law applies, yielding:

4. A is<sub>2</sub> not D.

Again, ‘B’ and ‘D’ refer to one thing, so, similarly,

5. A is<sub>2</sub> not B.

Line (1), by Leibniz’s law, makes (5) yield:

6. B is<sub>2</sub> not B.

By hypothesis, ‘is<sub>2</sub>’ expresses an identity, hence a reflexive relation. That, together with (6), is a contradiction. Call this argument *a*.

There is, though, another side. To begin, criteria are not always categorical, nor are categorical ones always available. Being male and being unmarried are criteria for being a bachelor: if you are to count as being a bachelor, then you must count as satisfying them. They sort categorically between bachelors and others only if there is always a unique right way of sorting people into those



who are married, or male, and those who are not. Suppose, e.g., a man has gone through a ceremony recognized in some places, but not in many others, as making him married. For some purposes we would say he was married; for others we would say he was not. As far as that criterion goes, such a man *might* also sometimes count as a bachelor and sometimes not.<sup>1</sup> Nor is there reason to think there are other criteria for being a bachelor which *are* categorical.

If there may be a man who, for some purposes, one could truly call a bachelor, while for others one could not, can we not also conceive of an object, *b*, in a similar situation with respect to its being *a* (though, for all that, there may be a criterion for being *a*)? The intuitive view, I think is that we in fact know of such situations. Here is one. In the seventeenth century, Vermeer painted a picture, 'Het Straatje', whose central subject is a house. Call it *Vermeer's house*. Around 1980, rubble was discovered in the cellar of a house. Call this *Kok's house*. Kok's house looked nothing like what Vermeer painted. The rubble, though, was shown to be the old façade of Vermeer's house. Kok's house, it seems, though different in nearly every respect, was the result of a series of renovations of (at the start at least) Vermeer's house. The intuitive view, I think, is that there are two views one could coherently take (though not both at once) as to whether Kok's house was Vermeer's house. One could view it as Vermeer's house, and would sometimes speak truth in saying so (as long as one's words were rightly understood as describing how things are on that view of the matter). Equally, one could view it as not Vermeer's house, and similarly sometimes speak truth in saying that (again, in words bearing a suitable understanding of what they thus say).

I said that this is the intuitive view. But one cannot sensibly view one object as another object; nor one object as two different ones. If Kok's house just is Vermeer's house, then any other view of the matter is inconsistent with the facts—that is, mistaken. Similarly, if Kok's house just is not Vermeer's house. In that case there cannot be two different *true* statements to make on the matter, one asserting, one denying of a certain house (being pointed out) that it is Vermeer's house. So if the argument set out above is correct, and if, as it seems to show, there is a categorical criterion for being Vermeer's house, then the intuitive view is wrong.

I think we should stick to the intuitive view. For one thing, we don't know how to give it up. (Which one thing is the right thing to say about Kok's house?) So I will argue that the objects we speak, and think, of do not need categorical criteria. To do that I will need to show how the above argument misconceives what is involved in the claim that pairs of identity statements may contrast in the way described, and why the intuitive view really is coherent. In doing that I

<sup>1</sup> There are those who hold that there could not, in principle, be a person who was anything other than married *punkt*, or unmarried *punkt*. I have argued against that view elsewhere. (See Travis 1994.) Even if it is correct, that is a large extra step beyond the idea that there are criteria for being a bachelor.

will try to identify just what picture of language and thought makes categorical criteria seem needed.

## II. Perspective

I will introduce an idea which I call perspective. To do so, I will assume, without argument, the following.<sup>2</sup> Consider (7):

7. The leaves are green.

Suppose (7) to be spoken of given leaves at a given time, so that 'are green' means what it does mean in English, hence speaks of being green. Suppose the leaves it speaks of had turned brown, and then been painted green—e.g., as a decoration. Then, so spoken—with the core semantics just described—those words might, and sometimes would, state truth, but also might, and sometimes would, state falsehood. They might do the former, e.g., where what is in question is whether the leaves are ready for the centrepiece, and the latter, e.g., where what is in question is whether the leaves are fresh.

The above means that, consistent with what it means, and while meaning that, the English predicate 'is green' may make any of many distinct contributions to wholes of which it is a part, and to the conditions for their truth. Distinguishing one of these contributions from others requires mentioning more than what is true throughout them—e.g., that 'is green' spoke of being green, was used to call something green, or describe it as being green. Each contribution requires the words to have a content they would have on some speakings but lack on others.

Why should different occurrences of 'is green' differ in what they contribute to saying of a thing? If in a particular speaking of (7), 'is green' makes some specific contribution to what was then said, rather than others it would sometimes make, compatibly with what those words mean, the circumstances of that speaking must decide that that contribution, and not any rival, is the one the predicate made that time.<sup>3</sup> Here is one description of how they do so. The English 'is green' is a means, in English, for calling things green.<sup>4</sup> So speak it of an item, and as meaning what it does, and you have called that item green, or done

<sup>2</sup> I have argued the point in, among others, Chs. 4 and 2 above and Travis 1989.

<sup>3</sup> One theory of *how* circumstances do this is that there is a specific set of parameters in occasions of speaking—illustrated by such things as time, place, and the identity of the speaker—and a specifiable function from values of those parameters to the content that the predicate would have on any occasion where those values held: as it were, a function—and I emphasize *function*—determined by meaning (or character, as it is sometimes called) from a specific domain of objects of some sort, where these objects are found in occasions of utterance, to content. It is at best obscure what the arguments, or even the values, of such a function might be. I will not pursue the issue here, but assume that this idea, in an area so far from the problems for which it was originally thought up (the variation in *reference* of indexicals like 'now' or 'here') has little future.

<sup>4</sup> Or describing them as green, saying them to be green, etc. Such fine distinctions will be ignored in this essay.

what then counted as calling it that. If so, you should have spoken truth only if the item (then) counted as being green. If that is how truth works, then, since what the predicate says varies across its occurrences, which items would count as being green must vary as well across occasions for raising those issues. Only if items which sometimes count as green may other times not are the various contributions of occurrences of 'is green' consistent with a constant connection between calling something green and being true only if it counted as *green*. The painted leaves, e.g., sometimes count as green and sometimes do not. That is *why* it is sometimes true, and sometimes false, to say them to be so.

If this is right, then it is wrong to think of the property, or concept, of being green as having such a thing as an extension, at least *tout court*. Since on different occasions, different things, and conditions of things, count as having *that* property, no function from objects to truth values can count as capturing correctly 'the' facts as to which items are green. Any such function must assign values to arguments which sometimes would, and sometimes would not, count as having the property in question. The best it could do while doing that is to capture which objects, in which conditions, would count, on some particular occasion, as being green. If it could even do that much, it would, in doing so, show no more about what it is to be green than would another function which did the same for another occasion. If, on one occasion, a concept counted as fitting some item, the way that item is, and on another, a concept counted as not fitting that item, while the item remains that way, that does *not* entail change of concept. Concepts and properties are not tied to extensions in that way; not, at least, if there is such a thing as the property, or the concept, of being green. If being green is a way for things to be, an item, while remaining the way it is, may sometimes count, and sometimes not, as being *that way*.

Precisely for that reason, some would deny that being green really is a property, or that there really is such a concept. *Real* properties, on this view, are not what is spoken of by English predicates, but rather what might be spoken of by the predicates of some more refined, perfected language. By definition, each has a definite extension *tout court*. Hence a predicate that expressed it would not exhibit the variation across speakings of it that the English 'is green' manifestly does.

It is at least obscure how, on this last view, we would ever arrive at real properties. For suppose we consider a 'pseudo-property' like being green, and a particular circumstance which might make an item, X, sometimes count, and sometimes not, as having that pseudo-property. The problem might be, for example, that paint counts, for some purposes, as making an item the colour it is, and for others as covering up its colour. Then we may introduce more refined 'properties' for which that issue is settled by stipulation—one, for example, on which if X, paint included, has that property, then X has it *tout court*. But is this more refined 'property' a property or a pseudo-property? Surely there are circumstances which would make an item sometimes count, and sometimes not,

as having it, too. Suppose, for example, that the paint was rather thinly applied, so that the result 'sort of' makes the dried-out leaf look green, but the brown does have somewhat of a tendency to show through. Or suppose the paint then becomes spotted with—say—bird droppings. So, it seems, more refinement is needed. Our new 'property' turns out to be a pseudo-property. It must be split into two, or four, or more, new 'properties'. When will we have achieved enough refinement? What we want are stipulations *so* explicit that there never could be an item which sometimes counted, and sometimes did not, as fitting them—as being that stipulated way. Our stipulations must decide in advance what to say about everything the world might ever conceivably confront us with. That, I suppose, over-taxes our powers of stipulation.

I will now use the above points to generalize radically on a point made by Frege:

If someone wants to say today what he expressed yesterday using the word 'today,' he will replace this word with 'yesterday.' (Frege 1918–19, p. 10)

To work out the semantic significance of this remark, consider a sentence such as

8. Today is cloudy.

Consider (8) spoken on day,  $D$ , so as to say  $D$  to be cloudy. Those are two of (8)'s semantic properties so spoken. But (8) has more semantics than that. There is more to the proper understanding of 'Today' in it. 'Today', so spoken, does not *just* refer to  $D$ . It does so in virtue of something else about it. Used on different days it would, *ceteris paribus*, refer to different days, where that is a consequence of what it means, and what, in virtue of that, any normal use of it represents itself as referring to. There is recognizably a feature of the content of 'today' with these effects. 'Today', in having that feature, identifies the day it speaks of (or represents itself as speaking of) as the day of its speaking. Call that feature  $F$ . As spoken on  $D$ , (8)'s semantics contains both  $F$  and the property of referring to  $D$ . Call that set of features  $S$ .

Now consider the next day,  $D + 1$ . Spoken on  $D + 1$ , (8) would not have  $S$ ; for, given that it retains  $F$ , it will lose the feature of referring to  $D$ . Instead, it will refer to day  $D + 1$ . It is not only (8) which, on  $D + 1$ , cannot have  $S$ . No words, on  $D + 1$ , could have that semantics. No words, that is, can, on  $D + 1$ , both present a day as the day of their speaking (have  $F$ ) and, in doing so, refer to  $D$ .<sup>5</sup> Call a set of semantic features a semantics. Then  $S$  is a semantics of a particular kind. It is available for some words to have on some occasions. Words produced on those occasions may coherently be supposed to have it. On other occasions, though, it is unavailable: no words could have it; we could not

<sup>5</sup> No doubt a *ceteris paribus* is needed here. Where people are badly misled as to what day it is, one of them might, perhaps, say 'today' and refer to a day other than the day he said it. There are also problems about some utterances of 'today' near the international date line, by telephone, etc. Such considerations do not spoil the present point.

coherently suppose that any then did. Let us call a semantics with this feature *perspectival*, and speak of words with it as expressing a perspectival thought. Call an occasion on which a perspectival thought would be expressed in speaking of thus and so (e.g., a leaf's being green) a perspective (on that).

Whether an occasion is one on which S is coherent is a matter of when it is. In other cases, the coherence of a semantics on an occasion may be determined by the place of a speaking. But it is not essential to the notion of perspective that we be able to specify precisely what features distinguish occasions on which a semantics is available, or coherent, from ones on which it is not. I now mean to abstract away from special cases where there is any such readily available means of saying what makes an occasion one on which a particular semantics is coherent.

Frege's idea generalizes if we take seriously the idea that there is a substantive task for circumstances to perform if words spoken in them are to make the sort of sense we in fact make of them. In given circumstances, within a given activity, or conversation, if you called certain leaves green, you would say what is true of them if they are painted green. If your 'are green' meant what it means, then you said that in calling those leaves green. Those are two features of the content of your words. Circumstances may confer that combination only by ruling out others. *Some* words, in calling leaves green, would say what is true only if the green is, as it were, natural. That pair of features is a coherent semantics for words to have. But it is not one words produced where, and in the way, yours were *could* have. That is not what it would be for a leaf to count, in your surroundings, as green.

Put more generally, the point is this. If, in saying 'It's green', you call something green, then what you thus say is true just in case that thing counts as green by the standards for being green which then apply. There are other ways of counting things as green or not, each of which is, on some occasion, the right way of doing it. Correspondingly, there are other semantics words could have while calling something green, besides the semantics your words did have. But those other semantics are not ones words could have where what counted as being green was what did count on the occasion of your speaking. The upshot is that words containing the predicate 'is/are green' typically express a perspectival thought, and have a perspectival semantics.

Some words 'They're green' are true if the leaves are painted green; some words 'They're green' are not. But there are no words 'They're green and they're green', the first half of which are true, given the paint, and the second half of which are false, given that there is only the paint, nor words 'They're green and they're not green' which are true, the first conjunct on grounds of the paint, and the second on grounds of what is under the paint. Conjuncts must be spoken from a perspective. If it is one in which paint counts as changing colour, then it is not one in which, in calling something green, you can say what is true, or false, despite the paint. So it is not one in which a conjunct of the envisaged sort is so much as producible. Similarly for other cases. There is literally no

saying such things. Variation in the content of 'is green' is essentially variation across perspectives. This exploits that aspect of the structure of thought which I signalled at the outset as making the initial intuition difficult to read.

### III. Perspectival Theses

Perspective enters into statements about being green in this way. There are many distinct things, any one of which may be said to be so of an object in calling it green. That is because there are various understandings of what it would be for the object to be green, each of them right in some surroundings for judging such matters. Where there is something the object's being green would rightly be understood to come to, in calling it green, one says it to be green on that understanding. One thus says one of the many things to be said in so calling it. Move through the space of available thought to another way of thinking about being green, or about that object's being green—move to circumstances in which there is a different thing its being green would come to—and in calling it green, one would state something else.

The thesis is that within a certain domain there are pairs of a true and a false statement, or true assertions and denials, each about a given object and the way it is at a given time. Consider the domain in question. The statements *need* not involve the word 'green'. For one thing, they need not be in English. But the domain should include all uses of the word 'green' with certain features. First, the word is used to ascribe some feature to an object, or to deny it to have some feature. Second, as used, it meant what it means (in English); so, as used, it says of the object whatever a word would say in virtue of meaning that.

Statements which share all the features just mentioned may yet diverge in the conditions under which they would be true. But they have something interesting in common. Suppose that, on some occasion, you wanted, while speaking correct English, to call some item green. You might say, 'It's . . .'. Then the word 'green' would be just the word for you to fill the blank with. So speaking, you will have done what then (in those circumstances) counted as calling an item green. That, and just that, is what the meaning of the word accomplishes, wherever it says what it does say in virtue of its meaning. The feature which connects all the statements in the domain is then this: each counted, on the occasion of its making, as calling an item green, or as denying that an item was green, and as saying no more nor less than that. It is unimportant whether the word 'green' was used for doing that.

Why is this domain interesting? Suppose we want to know what it is for an item to be green. We might examine what it is we speak of in speaking of being green. Now the point is: if that is meant to be fixed, or fixable, by what 'green', or a word which means what 'green' does, means, then each of the statements in this domain has equally as good a claim as any other to reveal what that is. So,

in fact, one may reasonably take what that is to be revealed only by what goes on throughout the whole domain. Each statement in the domain represents fully serious literal talk about whatever it is that 'green' speaks of. How else might we see what that is? The statements in the domain span many perspectives. No one of these distinguishes itself as *the* one from within which one can see what being green really comes to. Rather, what is so merely in virtue of having spoken of being green is what is so on any.

There is one alternative. That is the domain of all statements which say of an object what we would *now* say of it in calling it green. (This assumes what is probably not so: that for any suitable object, there *is* something we would now say to be so of it in calling it green. We are now doing philosophy; therefore more than likely to commit the sin, underscored by Austin, of trying to say, for no particular reason, what is and what isn't green, thereby saying *nothing* to be so. But bracket that point.) This domain is less interesting. Suppose that from our perspective there is a proper understanding of what it would be for an item to be green, which settles, for each item, what it would be for it to be green. That understanding yields a certain account of what being green is. But there are other perspectives, imposing other understandings. A philosopher enjoying one of these would, with equal right, come to a conflicting account of what being green is. Each of us is *right* about what counts as so from our own perspective. So far we do not conflict. Neither perspective shows better than the other what, independent of perspective, being green 'really' comes to. So neither account counts as a true account of that.

One way to bring out this point is to note how an account tied to one particular perspective would capture much information that should not be captured by a proper account. For example, from some perspectives, a leaf which is painted green thereby counts as green. An account developed from such a perspective would speak to the issue of whether painting a leaf green makes it count as green (as what we 'really' speak of in speaking of being green). But it is clearly no part of our notion of being green as such that paint has one sort of effect or the other. Our concept of green, as opposed to a concept of what would count as green on a particular occasion, just does not address the issues paint raises. Otherwise the contrasts of the last section would not arise. Green is a particular colour. To be green is to be that colour. Our understanding of what green is goes about that far and no farther. A proper account of what it is to be green ought to go no farther than that understanding allows for.

A question arises. If someone, or some words, counted on an occasion as then having called an item green, must he, or they, *now* count as having done that? The answer is complex. For expressions such as 'called green', 'said to be green', and even 'said that . . . was green' *may* bear either of two sorts of understandings. I will not pause to discuss the full issue here. There is, though, at least this much. If someone counted as having called X green, then one description of what he did is this: he said that X was green, on such-and-such

understanding of what it would be for X to be green. If so, then on *one* possible understanding of 'call green', he called X green. That is one fact 'called . . . green', meaning what it does, may state. Perhaps there is another understanding of 'call X green' on which to count as having called X green, one must have said X to be a way which *counts*, from the present perspective, as X's being green. On that understanding, if someone counts as having called X green, what he said is true iff X *counts* as (then) green, whether or not it so counted where he spoke. Counting as having called X green, on such an understanding, is not the mark of membership in the domain of interest. In any case, it is well to note that, given perspectival thought, the verb 'count' must do some serious work.

#### IV. Identity

The point about occasion-sensitivity was that for any concept there are many ways of applying it, different ones being correct in different circumstances. If that is true of *any* concept, it holds, *a fortiori*, of a concept of being such-and-such object. That would make a thought that *a* is a given object, *b*, perspectival, as thoughts about Vermeer's house and Kok's house seem to be. But what is the relevant thesis here? Its domain should parallel that of our thesis about being green. For the reasons given in that case, the statements in this domain have the best claim to be the ones to look at if we want to see what one says merely in virtue of speaking of the *identity* of some *a* and some *b*. The variety in this domain, the differences between one thing and another to be said within it, shows what the features are which distinguish one understanding of an identity statement from others; what content such a statement has.

We may define the relevant domain thus. A statement is in the domain if, for some *a* we may speak of, and some *b* we may speak of, the statement counted, at its making, as speaking of *a*, as speaking of *b*, as speaking of their identity (their being one, or being two), and either as asserting, or as denying precisely that. Such an assertion of identity will count as true just in case the objects it then counted as having spoken of then counted as one; such a denial just in case they then counted as two. The possibility is left open that words may have counted, at their speaking, as having spoken of a certain *a*, or of identity, without now counting as having done that. Whether, e.g., a statement states *identity* may be a matter of perspective.

The thesis that identity statements express perspectival thoughts concerns this domain. It is that there are pairs of members of the domain, one true and one false, where, for some *a* and some *b*, each counted, at its making, as speaking of *a*, speaking of *b*, and asserting their identity; and there are pairs in the domain, both true, each as before, except one having counted as an assertion, the other as a denial, of the identity of the objects they speak of.



On the perspectival thesis for green, there may be *something* true to be said in calling an item green without it *now* being true to say that that item is green. That true thing need not be what would, or could, be said from the present perspective. With that in mind, we can see the perspectival thesis for identity not to be threatened by argument *a*. For on the thesis, if there is a true statement identifying *a* and *b*, that is what one would say in speaking of their identity from a particular perspective: namely, one from which what counts as being *a* and what counts as being *b* is such that what counts as *a* also counts as *b*. There is simply no truth to be told from such a perspective in denying that *a* is *b*. If there is, nevertheless, *something* true to be said in such a denial, that can only be in a denial made from another perspective. It is not something that could be said at all from this one. There is thus no perspective from which both premisses of *a*, so *a* itself, could be so much as stated. What could not be stated is not genuinely an argument. In mixing perspectives as it does, *a* is no more than an ill-formed would-be argument.

But we are not yet home free. We want to be able to say such things as this. Jones pointed at Kok's house, said 'This is Vermeer's house', and spoke truth. Smith, on a different occasion, pointed at Kok's house, said 'This is Vermeer's house', and spoke falsehood. Jones's 'This' counted, when he spoke, as speaking of Kok's house. Smith's also so counted as he spoke. Jones's 'Vermeer's house' similarly counted as speaking of Vermeer's house. Similarly for Smith's 'Vermeer's house'. There is now this objection. Now—as I say all this about Smith and Jones—either Kok's house counts as Vermeer's house or it does not.<sup>6</sup> Suppose it does so count. Smith's 'This' counted as speaking of Kok's house. Since, by hypothesis, Kok's house is Vermeer's house (we now state truth in saying so), that 'This' counted as referring to Vermeer's house. Since it thus counted as referring to that very thing to which Smith's 'Vermeer's house' counted as referring, what Smith said, if an identity statement, cannot be false. Things are equally bad for our description of Jones if, from our own perspective, Kok's house does not count as Vermeer's house.

If the perspectival thesis is correct, the following form of inference is invalid: 'N', on a given speaking, counted from such-and-such perspective as speaking of *a*; *a* is *b*; so 'N' then counted as speaking of *b*. *Should* this be invalid? Let *a* be a house, and *b* be a house. There are, first, many houses to speak of; many any given house might be or not. There is, e.g., Vermeer's house. There is, e.g., Kok's house. I have just used two means for speaking of a house, each available in suitable situations. In making use of each, I, and my words, represent ourselves as identifying a particular house. If people, or words, ever succeed at such things, then I, and mine, did. These same means of identifying a house, these same ways of saying which house it is, are available on other occasions. Just as for

<sup>6</sup> I do not want to introduce a third possibility here, since no such move is needed to make the perspectival thesis coherent.

any concept, or description of things, there are many ways of applying it, so for any means of speaking of a house there are many ways of applying that means; correspondingly, many different standards, each sometimes right, for judging what, if anything, one thus speaks of or identifies. Given that much, the inference in question ought to fail. For to count, from some other perspective, as speaking of *a*, or as speaking of *b*—to do what *then* counts as that—is to do so by what are *then* the right standards for judging such things. Many ways of deploying a given identification implies many standards, each sometimes right, for judging what, if anything, is thus identified. Once perspective is in our picture of thought, we cannot make something count in other circumstances as so just because by present standards for judging things it counts as so. If *a* is *b* and the person now counts as having spoken of *a*, then he *now* counts as having spoken of *b*. Our *present* standards for judging take us no further.

It might be objected: if there is something that would have counted as an object's being the one 'A', then spoken, spoke of, but would not now count as an object's being the one 'A', now spoken, spoke of, then these two 'A's could not be the *same* way of identifying an object; nor could there be an *a* which each counts, or ever counted, as identifying. That idea of how to speak of objects misrepresents our capacities to think of objects—what we could grasp in grasping a way of identifying an object and in grasping it as a way of doing that. Compare green. In speaking of being green I identify a way for an object to be: green. In saying an object to be *that* way, one might say any of many distinct things of it. *That* way for things to be does not decide, in a unique way, what would count as a thing's being that way. No perspective-independent facts as to what *is* that way make it the way it is. Nor would this picture change for some future 'more precise' concepts. However much we grasp, or could grasp, of how a thing must be to be a given way, there are always various understandings of when a thing will have done all it thus must. That leaves various things that might sometimes count as speaking of that very way for things to be. Similarly, no matter how much we may grasp of what an object must be to be the object thus and so, that allows various understandings of when an object would be that, hence various things that might sometimes count as speaking of that very object. Such is what one would expect of the capacities of humans, as opposed to gods. Without the non-compulsory inference pattern above, nothing in logic forces a more hubristic view of our capacities to fix what it is we speak of.

Frege held that we need a sense, or mode of presentation, as an element in the proper understanding of a name in order to identify what statements of identity say, distinguishing one such thing from another. Such statements speak, to be sure, of a relation between objects; one holding precisely between each object and itself. The rub comes in distinguishing the different things one might state, the different ways one may say things to be, in speaking of *that* relation. If a sense is a path to a referent, being at the terminus of one such path is a different accomplishment from being at the terminus of another, even if some one object

achieves both things. An identity statement credits an object, at the end of one such path, with another such achievement. Difference in sense, in such a context, thus results in saying different things to be so.<sup>7</sup>

I have spoken of means of identifying an object. ‘Vermeer’s house’, understood as speaking of, or expressing a concept of being, a certain house, namely, Vermeer’s, is an example. Such identifications play a role here similar to Frege’s senses. I agree with Frege that identity statements speak of a relation between *objects*. I also agree that we must recognize an extra element—here the means used to identify the object spoken of—in order to say what it is that is so according to identity statements. For, again, counting as what is identified by one means and counting as what is identified by another are two different achievements. Crediting objects with different such achievements is to say different things about them, though, for all that, it is to speak of *them*.

Frege also held that *no one* sense can, on different occasions, refer to different objects. Nor have I said that an identification, or means of identifying, may, on different occasions, pick out different objects. The point is not that a true ‘Vermeer’s house is Kok’s house’ and a false one differ in that, in each, ‘Vermeer’s house’ refers to a different object. In each case it counted at its speaking as referring to Vermeer’s house. Each occurrence might well now count as doing that as well. It is not as if there is some *other* house for them to refer to, nor as if one, at least, spoke of something other than a house. But here I adumbrate topics in the next section.

Wherein, then, does the present account depart from Frege? Frege thought, not just that each sense must in fact have just one reference, but that a sense is consistent with precisely one set of semantic properties full stop—in this case, just one set of facts as to what would, and would not, count as what a name with it named. A sense leaves no room for a *variety* of understandings of what it would be to be what words with it speak of, or of how they say things to be. So a sense for ‘Vermeer’s house’ leaves no room for Kok’s house sometimes to count, and sometimes not, as what it names. Precisely that is denied by the present account. The idea is: no understanding *we* can give words precludes a variety of possible understandings of what it is they would say, or speak of, on that one; no semantics, or sense, *our* words can bear is fit to be embedded in but one larger set of semantic properties. We cannot understand an identification of an object so that, so understood, it is fit to make only one contribution to any

<sup>7</sup> If we introduce a notion like sense, we can make sense of the existence of pairs of statements related in a certain way. But for Frege it was precisely the existence of such pairs which was supposed to demonstrate that there were senses. So, one might think, there are senses if there are senses. Things are not that bad. What we must look at is the reason we would have for recognizing such pairs if we could do so coherently. A theory of content with a notion working, in this respect, like sense, may recognize perspectival identity thoughts and statements. In so doing, it accounts for a host of intuitions which a theory without such a notion must dismiss as bogus. If there is good reason to accept these intuitions as genuine manifestations of what the competent are prepared to recognize, that speaks in favour of recognizing such an element in understandings.

whole statement of which it is a part, nor get an identification to bear such an understanding.

There is a picture on which if we do, through our understandings, latch on to an object, then what is and is not that object has nothing further to do with *understandings*, any more than whether an utterance was loud depends on what understanding it bears. Objects have their histories quite independent of us. Of course they do. So does the colour green—first colouring this, and then that. For all that, there are various understandings of when an object would be that colour, and that does matter to when an object so counts. If objects have their histories, what counts as part of *one* such history? Here the possibility arises for a variety of views, and with that a role for understandings, parallel to that in the case of green, arises once again.

## V. Objects

There is no incoherence in the idea of pairs of identity statements related as the perspectival thesis has it. Neither logic nor Leibniz's law prohibits this. So if, by our ordinary standards for speaking truth, someone might, on an occasion, point at Kok's house and say 'This is Vermeer's house' and speak truth, using 'Vermeer's house' to speak of Vermeer's house, and, also, someone might, on a different occasion, point at Kok's house and say 'This is not Vermeer's house', using 'Vermeer's house' in that same way, and also speak truth, then the perspectival thesis is true. It is also true if there is any other similarly related pair of things to say. We may now confirm that if the perspectival thesis is true, then the Fregean view of the relation of reference to truth is wrong. For the sake of argument, let us suppose that the case of Vermeer's house actually works as suggested.

On one occasion, Jones said 'This is Vermeer's house' and spoke truth. On another Smith said 'This is not Vermeer's house' and spoke truth. Each pointed at Kok's house, and each looked the same while doing it, as did Kok's house. *We* now want to say what each referred to. A natural description is this. Jones, in saying 'This', referred to Kok's house. In saying 'Vermeer's house', he referred to Vermeer's house. Smith did likewise. If that description is correct, then Frege's view of the relation of reference to truth is wrong.<sup>8</sup> For we have correctly said Jones and Smith to have referred, piecewise, to the same things. So we have correctly represented their statements as agreeing in reference. But, counter to what Frege's view predicts, both are true. (Smith's statement is thus the negation of a falsity agreeing piecewise in reference with Jones's truth.)

To preserve Frege's view, we need to postulate differences in reference between Jones's statement and Smith's. The most straightforward way to do so would

<sup>8</sup> Given a reasonable view of negation.

be to say roughly this: Jones's 'Vermeer's house' referred to an object which just is, *tout court*, the sort of thing that survives that sort of rebuilding. Smith's 'Vermeer's house' referred to a different object, whose nature is such that it does not survive such rebuilding, *tout court*. Jones's object was a house, perhaps (if that is how houses behave). Smith's object must then have been, as one might call it, a shmouse. (It is not, e.g., a mere part of the house, such as the façade; nor is it something then at some other address.) Houses and shmouses, along with myriad other things, tend to, as one might put it, cohabit: for the life of the shmouse it existed at the same place, and was made of the same bits, as the house. It's just that at a certain moment the shmouse gave up the ghost, while the house went on.

Without perspective we would be driven to proliferate objects in this way. For we said that Jones's 'This' and Smith's 'This' both referred to Kok's house. And Jones's 'Vermeer's house' and Smith's both referred to Vermeer's house. (All that is what counts as so from our perspective.) Without perspective, either Kok's house is Vermeer's house or it isn't, full stop. (The one thing or the other, but not both, is what it is true for us to say. That much is so with or without perspective.) Without perspective, in the first case, Smith cannot have stated truth, and in the second, Jones cannot have. (If it is possible for Kok's house and Vermeer's house to count neither definitely as one nor definitely as two, that would only mean that neither Jones's statement nor Smith's could be true.) So if we were right about the truth values of Jones's words and Smith's, then there must be differences in reference of the sort just indicated. (More ordinary differences, such as referring to the window rather than the door are ruled out by what was in common to what each did.)

Perspective means, though, that we are *not* driven to such remedies. For, given it, the truth of, e.g., Jones's statement depends, not on whether from *our* perspective Kok's house and Vermeer's house are one, but rather on whether in doing what then counted as speaking of Kok's house (and, as it happens, doing something which also so counts from our perspective), and, similarly, in doing what then counted as speaking of Vermeer's house, Jones did something which *then* counted as speaking of one object twice. Given perspective he might have done so whether or not one could perform a similar feat from our current perspective.

The perspectival thesis thus allows us to say what we should anyway want to say about reference to objects, and to reject proliferating objects as above. What we should say is that such proliferation depends on a wrong view of what it is for reference to be determinate. On a proper view, at least in normal talk, speaking of shmouses rather than houses is not something it is open to us to do. There are not normally two distinct references to be made, distinguished from another in that one is to a house, the other to a shmouse. Shmouses are not normally among the objects of which we might speak at all. If, e.g., Jones points to Kok's house and says 'This house is old', what he purports to refer to is a *house*; and he does

so if anyone ever does. The problem is then to say which house. Among all the houses there are, there is no better candidate than Kok's. That is how answers to such questions go. And such an answer does no less than make it determinate which object was referred to. That is what being determinate looks like.

There are now two things to say. First, if someone counts as having referred to the object such-and-such, or the one with such-and-such features, and if that object, or the one with those features, counts as having such-and-such other feature, then the person counts as having referred to an object with that further feature. That is a rule to be applied from within any perspective. So if Smith now counts as having referred to a certain house—say, Kok's—if that counts as answering the question what Smith referred to, and if that house now counts as dating from the seventeenth century, then Smith now counts as having referred to something dating from the seventeenth century. Smith cannot count as referring to a house, but not a seventeenth-century one, where the only house for him to have referred to counts as a seventeenth-century one. Nor is there any suitable non-seventeenth-century object for Smith to have referred to, as there would be if there were shmouses, and if he might have counted as having referred to one.

Second, facts which make it determinate which object is in question may leave it open which further features that object counts as having, so that the object made determinate in that way may, for all that, count, from some perspectives, as having features which, from others, it counts as lacking, and may do that even though, of course, it is, throughout, the way it is. Such features may include, of course, such things as being green. But they may also include such things as surviving, or pre-dating, certain operations or happenings, and so too such features as having existed in the seventeenth century or still existing now. This entails that among such features may also be a feature of being such-and-such object, the object A, where these last words speak of an object made determinate in some other particular way.

Given perspective, then, one cannot save Frege's view that the references of parts fix a unique truth value for their whole, in the case of identity statements, through differences in which objects are referred to in suitably chosen contrasting pairs. For, normally, there are not both houses and shmouses to be referred to. And in a well-chosen case that is the sort of difference there would have to be. Given the broad notion of reference in play in Frege's view, one might hope, in such cases, to find differences in the reference of some predicate. In each statement, e.g., there is a predicate which speaks, not just of Vermeer's house, but of *being* Vermeer's house. Can we not find *differences* in what each predicate (instance) speaks of?

Perhaps so. One speaks of being Vermeer's house on *this* understanding of what that would come to; the other of being Vermeer's house on *that* one. That merely shows that we *can* describe these two occurrences of the predicate 'is Vermeer's house' as speaking of different things. It remains so that there is something such that we can describe both of them then as speaking of *that*:

namely, being Vermeer's house. Do that, and similar things, and you *have* said what the parts speak of. But you have not said what fixes a unique truth value.

Moreover, if there are other ways of saying what a whole's parts speak of, what the parts are thus said to speak of may still allow for more than one truth value for wholes which speak of that. Suppose we say that a given instance of 'is Vermeer's house', in a given identity statement, speaks of being that house on such-and-such understanding of the matter. There may still be various things to be said of an object in saying it to be Vermeer's house on that understanding; ways for an object to be which, from some perspectives would, and from others would not, answer to that understanding of what Vermeer's house is. This repeats a now-familiar point about the supposed progression, through successive refinements, towards properties which have, as such, a definite extension. In this context the point becomes: we are still not *en route* to a sense in which referents of parts determine a unique truth value.

Leibniz's law and its consequences, transitivity, commutativity, and reflexivity—like any principles governing inference—hold from within a perspective, though from within all. *Whenever* I count *a* as *b*, I must always *then* accept all the consequences that Leibniz's law, applied within that perspective, yields. *Healthy* realism must also be what is expressed from a perspective. I cannot change one object into two, or two into one, by, as it were, smooth talk (or thought). An object stands there, as one might say, like a cow. It is just the way it is. Perspective does not challenge that.<sup>9</sup> Nor does any of that rule out, as we move across perspectives, finding, from one, one object encountered twice where, from another there are two objects. *Healthy* realism must live with that.

<sup>9</sup> It is also compatible with the principle that if *a* is *b*, then necessarily *a* is *b*, which again is a principle governing talk and thought from within a perspective.

# 12

## On Constraints of Generality

*Waismann:* The very possibility of concept-formation shows that every proposition belongs to a logical space which is nothing but the system.

*Wittgenstein:* If ' $\phi a$ ' is supposed to be a proposition, then there must also be a proposition ' $\phi b$ ', that is, the arguments of ' $\phi()$ ' form a system. What I admittedly do not know is how large the domain of arguments is. And there might, for example, be only two. . . .

But does ' $\phi a$ ' presuppose ' $\psi a$ ' too? Decidedly yes. For the same consideration tells us: if there were only a single function ' $\phi$ ' for ' $a$ ', then it would be superfluous; you could leave it out. The propositional sign would be simple and not composite. . . .

Result: a proposition can be varied in as many dimensions as there are constants occurring in it. The space to which the proposition belongs has just as many dimensions.

A proposition reaches through the whole of logical space. Otherwise negation would be unintelligible. (Waismann 1979, pp. 90–1)<sup>1</sup>

### I

Gareth Evans proposed, as a constraint on thought, that anyone, to entertain any given thought, must be able to entertain a particular (thinker-relative) system of them. Specifically,

We thus see the thought that  $a$  is F as lying at the intersection of two series of thoughts: on the one hand, the series of thoughts that  $a$  is F, that  $b$  is F, that  $c$  is F, . . . , and, on the other hand, the series of thoughts that  $a$  is F, that  $a$  is G, that  $a$  is H, . . .

. . . if a subject can be credited with the thought that  $a$  is F, then he must have the conceptual resources for entertaining the thought that  $a$  is G, for every property of being G of which he has a conception. This is the condition that I call 'The Generality Constraint'. (Evans 1982, p. 104)

Evans states the constraint for the second series, but clearly takes it to entail, as it does, a result about the first. Suppose I can entertain, or grasp, a thought about

<sup>1</sup> I thank my colleague Peter Sullivan for calling this passage to my attention.



some item to the effect that it is G. Then I know what being G is. Suppose I can think *some* thought about *a*—say, a thought that *a* is F. Then, by the generality constraint, I can think, or grasp, or entertain, the thought that *a* is G. So, if I can think any thought to the effect that something is G, then I can entertain, or grasp, the thought that *a* is G, for every item *a* which I can think about at all. It is this version of the constraint that Evans works hardest.

**Note 1:** Evans usually speaks of ‘entertaining’ thoughts. This, I take it, is to deal with the fact that to think a thought is, *inter alia*, to believe it, not just for some inner voice to mouth it. I will usually speak of grasping thoughts, rather than of entertaining them. One grasps a thought, I take it, in knowing which thought it is; in being able, adequately, to distinguish it from other thoughts. One does that, I take it, if one grasps adequately what it would be for things to be as they are according to that thought, or, more simply, what it is that is so according to it.

**Note 2:** Evans recognizes that the second version of his constraint (the version for the first series) may need modification on account of facts about categories. Perhaps we can grasp the thought that Dobbin eats hay without being able to grasp the thought (if there is one) that the number 2 eats hay. Nothing I have to say exploits the possibility of category mistakes. I will consider *only* predications which are no such mistake; thoughts of what might, in fact, be so.

I will argue that our thinking is not constrained by the generality constraint. So there is no legitimate philosophical work that constraint can do. The constraint seems plausible only given an over-simple picture of which capacities are in play in grasp of a thought. Insofar as our grasp of a particular thought to the effect that *a* is F involves grasp of what it is for a thing to be F, it also requires exercise of a further sort of capacity which Evans ignores, and which may be present with respect to a given thought to the effect that *a* is F, while missing for another thought to the effect that *b* is F (where *a* and *b* may, or may not, be distinct).

Nothing I will say conflicts with Wittgenstein’s idea about systems. For Wittgenstein, a system of thoughts has certain roles to play in conferring the content they have on each of its members. It may well be that systems are needed to play such roles, and hence that any thought must be understood in terms of some system of thoughts to which it belongs. Perhaps what I have to say will even suggest further reasons for thinking that—as Wittgenstein developed further reasons for thinking it in his later philosophy (esp. 1969). If the plausibility of the generality constraint derives in part from that of Wittgenstein’s idea, it is important that the two ideas are by no means the same.

My main point is equivalent to one Wittgenstein makes as follows:

You say to me: ‘You understand this expression, don’t you? Well then—I am using it in the sense you are familiar with.’—As if the sense were an atmosphere accompanying the word, which it carried with it into every kind of application.

If, for example, someone says that the sentence ‘This is here’ (saying which he points to an object in front of him) makes sense to him, then he should ask himself in what

special circumstances this sentence is actually used. There it does make sense. (1953, §117)

Another version of the point is:

Surely  $\sqrt{-1}$  must mean just the same in relation to  $-1$  as  $\sqrt{1}$  means in relation to  $1$ ! This means nothing at all. (1953, p. 190)

Wittgenstein's point is directed against the philosopher who thinks that what is said, and what is said to be so, in any occurrence, or speaking, of words—the thought they express—is fixed, modulo a small number of quite regular factors (such as time and place referred to), by what the words *mean* in their language (assuming that they are well formed, and ignoring the possibility that they may bear several readings in their language). Whereas, the point is, what words mean (the aura they always carry with them) is compatible with their saying an indefinite variety of different things (of given objects, times, and places). Evans supposes, in effect, that if we fix as much about a thought as what words mean, plus their referents, fix about which thought they expressed, then we have fixed which thought is in question. If one grasps that much about a thought, then one grasps the thought. That, I will argue, is not so. There is another identifiable capacity which must be exercised in conjunction with knowledge of what F is before that knowledge can yield a grasp of what it is that is so according to any particular thought to the effect that *a* is F. That capacity may be missing for any given thought that such-and-such is F, even given the ability to think about that object plus a grasp of what being F is adequate for grasping some other thoughts that such-and-such is F. Hence the generality constraint cannot hold.

## II

One intuitive idea runs counter to the generality constraint. Suppose, for example, someone says that the leaves on the tree are green. Fine. We understand what it would be for things to be that way; we grasp the thought expressed. Now suppose someone says that his bedroom walls are green. Again, we grasp that thought; know how things would be according to it. If someone says that the cheese we left in the refrigerator when we went on vacation is green, again, so far, so good. Now suppose someone calls his Uncle Hugo green. Might we not, for all of the above, be baffled as to what is supposed to be so according to that thought, unhelped by our knowledge of what being green is, adequate though it was for grasping those other thoughts? Would we not, for all that, know what being green is?

Someone might think this intuition compatible with the constraint. The thought would go: the intuition points to a sort of possibility of failure to understand *words*, whereas the constraint is meant to concern thoughts. If someone produces given words, such as 'Uncle Hugo is green', then one of

the problems that would be solved in, or by, understanding them would be associating the right thought with the words; discerning, that is, which thought they actually expressed—this one, rather than that one. If there are several thoughts that might be expressed in calling Uncle Hugo green, one might fail to know which one was expressed in the words in question. That still does not show that there is any particular such thought which one would be unable to grasp—once one supposed given words to express it.

One of my tasks in what follows will be to dispel the idea that the problem the intuition points to is one merely about words, disappearing at the level of thoughts. The idea that thoughts may eliminate, in principle, problems which inevitably arise for words is anyway suspicious. Wittgenstein referred to thoughts, conceived as fit to do that (providing interpretations of words, without themselves admitting of interpretations) as shadows, and said:

even if there were such a shadow it would be susceptible of different interpretations just as the expression is. . . . You can't give any picture which can't be misinterpreted. . . . No interpolation between a sign and its fulfilment does away with the sign.<sup>2</sup>

I will argue that Wittgenstein was right.

On the intuition, it may be one thing for *a* to be F, and quite another for *b* to be F; different considerations may be relevant to deciding the correctness of different predications of F of something. That idea is reminiscent of Wittgenstein's discussion of family resemblance (1953, roughly §§65–92). The similarity is more than coincidental. Family resemblance is one of the devices Wittgenstein uses in about the first 120 paragraphs of the *Investigations* to set out his most central ideas on words and their relation to thought. Those ideas clearly show why it is wrong to insist on the generality constraint.<sup>3</sup> Let us consider briefly how the opposition between family resemblance and the generality constraint works.

One main idea of family resemblance can be put this way. Consider a given concept; one of being F. Now consider its application to some object, *a*. Think of all the ways for *a* to be which would count as *a*'s fitting the concept (being F), and all those which would not; of the sorts of considerations in virtue of which *a* would, or would not, so count. Now consider a new application of F—say, to some other object, *b*. In the same way, survey what would, and what would not, count as *b*'s being F. One may detect similarities and overlaps, but also differences in the considerations which matter in each case—those, e.g., that would disqualify *b* from counting as fitting the concept and those that would disqualify *a*. Surveying more occasions for predicating F of a thing, we see that there is no one set of considerations which is always the set which decides whether such a predication is correct, distinguishing what would count as some thing's being as thus said to be from what would not.

<sup>2</sup> Quoted by G. E. Moore in his 1993/1954.

<sup>3</sup> It is notable that in discussing family resemblance Wittgenstein is thoroughly indifferent to whether his points are put in terms of words (predicates) or concepts. That is a straw in the wind.

The point is not that there are no features in common to all cases in which something counts as F. That, for Wittgenstein, would be nonsense. After all, all things which count as F share this in common: they count as F. Nor need Wittgenstein deny that, in particular cases, there are more features in common to all things which count as F. The key point is better put the other way around. In the case of a particular application of F to an object, *a*, there will be facts such that, if they obtain, then *a* does not (on that application) count as F, and, perhaps, facts such that, if they obtain, *a* does so count; which facts have such status in particular cases is not derivable merely from any correct account of what is in common to all cases in which something counts as fitting that concept.

The idea here comes out in Wittgenstein's example of the radical. There is something in common to any pair of numbers which stand in the relation *x is a square root of y*: namely, that the first is the square root of the second. One can say more: the first is such that the result of multiplying it by itself is the second. We also grasp what is required for any number to bear that relation to 1, and grasp this in a way which allows us to say which numbers do, and which numbers do not, do that. Now ask what is required for a number to stand in that relation to  $-1$ . Which numbers do, and which do not, bear that relation to  $-1$ ? Nothing mentioned so far answers that question. Whether there are any square roots of  $-1$ , and if so, what they are, depends on the kind of arithmetic we choose to do. If there are facts which decide, for particular applications of the concept of the square-root relation to pairs  $\langle a, -1 \rangle$ , whether, on those applications, that pair counts as standing in the relation, those facts do not follow merely from any general account of what the square-root relation is, nor from that plus as many facts as you like as to what is required for something to bear that relation to 1.

The generality constraint conflicts with the idea of family resemblance in supposing that there is some general understanding of what it is to be F which, for any object *a* one can think about, is enough to yield an adequate understanding of what it is for *a* to be F. In that, I will argue, it is mistaken.

### III

As Evans sets things up, grasp of 'the' thought that *a* is F plus grasp of some (any) thought about *b* jointly entail the ability to grasp 'the' thought that *b* is F. What is needed to grasp a thought presumably depends on which features distinguish it from other thoughts, true on, or of, different conditions. It might thus be well to ask how many thoughts there are, all about *a*, and about being F, and to the effect that *a* is that, each true on, or under, different conditions.

There is one traditional idea on which the answer to this question is 'many'. I mention it only to set it aside. Since Frege, it is often held there may be two thoughts to the effect that *a* is F, distinguished from one another by their 'mode

of presentation' of *a*, that is, I suppose, by the way in which, in thinking, or entertaining that thought, one would be thinking of, or identifying, *a*. Though something like that may be right, I am not interested here in that sort of variation. What I want to examine are the different contributions a given concept of being F may make to different thoughts to the effect that a given object (which might as well be thought about in a given way throughout) is F. I will speak of different thoughts only where there is a way things might be, or might have been, such that if things are (were) that way, then things would be as they are according to the one thought, but not as they are according to the other; so that there is a way for things to be on which the one thought would be true, the other not.

Although the point I am after is about thoughts, I begin with one about words. As an arbitrary example, consider the words 'The leaf is green', speaking of a given leaf, and its condition at a given time, used so as to mean what they do mean in English. How many distinct things might be said in words with all that true of them? Many. That emerges when we note that one might speak either truth or falsity in such words, if the leaf is the right way. Suppose a Japanese maple leaf, turned brown, was painted green for a decoration. In sorting leaves by colour, one might truly call this one green. In describing leaves to help identify their species, it might, for all the paint, be false to call it that. So words may have all the stipulated features while saying something true, but also while saying something false. Nothing about what it is to be green decides whether the colour of a thing is the way it is with, or the way it is without, the paint. What being green is, is compatible with speaking either truth or falsity in calling the leaf green. For all that, the painted leaf is as it is sometimes, but not other times, said to be in calling it green.<sup>4</sup>

The words 'is green', while speaking of being green, may make any of indefinitely many distinct contributions to what is said in words of which they are a part. The above variation is illustrative. The same holds of any English predicate. The fact that 'is green' speaks of being green does not alone decide what is required for a thing to be as it, on a speaking, says a thing to be. Similarly for whatever else words speak of.

On minimal assumptions, the point applies immediately to thoughts. Suppose we use 'thought' so that words which say what is true under different conditions *ipso facto* express different thoughts; the *same* thought cannot be sometimes true on one condition, sometimes true on another. Then words, while speaking of a given leaf and calling it green, may express any of many thoughts. So, equally, a thought may be about the leaf, about its being green, and to the effect that it is that, and for all that be any of many thoughts.

There are now difficulties about what the generality constraint is supposed to say. Is the grasp of *one* thought to the effect that *a* is F enough to guarantee the

<sup>4</sup> I do not plan to defend this thesis any further here. <sup>4</sup> For further discussion, see Chs. 2 and 4 above, or my 1989, esp. ch. 1.

ability to grasp every thought, of anything, to the effect that it is F? Or is it only enough to guarantee the ability, for every  $b$ , to grasp some such thought? Or etc.?

Or perhaps the constraint is meant to be something else. How is it, one might ask, that the form of words 'the thought that  $a$  is F' has a perfectly good use, despite the fact that there are many thoughts equally to the effect that  $a$  is F? How do those words manage to refer to a particular thought? The answer is this. Though 'is green' may make any of many contributions to a whole of which it is a part, I may use it on an occasion so that it makes some one contribution: in some surroundings for its use there is some one contribution it would make—one on which there is some particular thought expressed. If it is so used in the expression 'the thought that the leaf is green', the result will be an expression which refers to that thought. Now perhaps the idea of the generality constraint is this. Someone who grasps the thought that the words 'the thought that  $a$  is F' refer to, spoken on a given occasion, or where 'is F' makes some particular contribution from the many it might, is thereby able to understand the thought that would be referred to by 'the thought that  $b$  is F', spoken on that occasion, or where 'is F' makes the same contribution. This presupposes that, e.g., it is determinate what it is for 'is green' to make the same contribution to the expression of a thought about a leaf as it does to a thought about mouldy cheese, or Uncle Hugo, and/or that surroundings which determine what thought would be expressed in calling a leaf green must, inevitably, determine what thought would be expressed in calling Uncle Hugo green, and that so it is with predicates in general. In any event, no such version of the constraint is correct.

#### IV

For any  $a$  and F, there are many thoughts, each thinkable on some occasion or other, all equally about  $a$ , and about its being F, each differing from the others in when things would be as they are according to it. What capacities would be involved in grasping some one of these, and the way things are according to it, as opposed to other ways things are according to thoughts distinct from it? Again, I will seek an answer by looking first at understanding words. I will then plead the usual Wittgensteinian point: thoughts do not escape what words in principle could not. Problems of grasping what is so according to an item are not eradicated by making that item a thought instead of words.

Consider some entirely typical words: the predicate 'is green'. There are various things to be said to be so of a given object  $a$ , in saying ' $a$  is green', using 'green' so as to speak of being green. That has an immediate consequence: what 'is green' means cannot by itself determine which contribution those words make to what was said in using them on an occasion. So the circumstances of that use, or as I shall say, the surroundings of that speaking, must contribute substantially to determining which of various possible things that contribution comes to.

Parenthetically, someone might hold that what ‘is green’ means determines a set of parameters (variables in speakings), and a function from values of them on to a range of contributions ‘is green’ might make, such that for any argument of the function (fixed relevant values of the speaking), the value of the function is that contribution which ‘is green’ would make on a speaking so characterized. The model for this idea is the supposed way in which what ‘now’ means fixes a parameter—the time of utterance—and a function—the identity function—such that ‘now’ on any speaking refers to the value of that function for the time of speaking as argument. Even if that were so for ‘now’, there is no reason to think that what ‘green’ means similarly determines any such set of parameters, or that there are any. I will henceforth suppose that there are not.

How might surroundings do their work? There is a framework for an answer in the idea that meaning is use. There is something the predicate ‘is green’ is for in English: a use it has, and a use to give it in speaking English correctly. For one thing, to oversimplify, it is for calling a thing, or describing it as, green. Use it in speaking English, so that it is to be understood as doing what it is for, and that is what you and it will have done.<sup>5</sup> For another, still to oversimplify, it is for describing green things. If you aim to use it correctly, so as to describe what it does describe in English, then—at first approximation—those are the things you ought to use it to describe.

If you call a leaf green, the leaf must count as green for what you thus said to be true. But *when* must it so count? There is more than one thing to be said in calling a leaf green. It may be that the way the leaf is makes some things thus to be said true, and others false. So it cannot be that there is a fact as to the way the leaf is—it counts as green, or it doesn’t—such that that fact decides the truth or falsity of all such things to be said. Rather, the truth of different such things is decided by different considerations about the way things are. But the initial intuition is still right if read right. The right reading is this: if, in given surroundings, you call a leaf green, what you thus say is true just in case the way that leaf is counts, in those surroundings, as its being green. If, given the way the leaf is, you sometimes speak truth in calling it green, and sometimes do not, that can only be because in different surroundings different things would count as the leaf’s being green. A leaf may count as green *on a particular understanding* of what its being green would come to; circumstances may make one understanding or another the right one.

Why is it that what would count, in some surroundings, as a leaf’s being green would not do so in others? An answer can be: not all of what words are for is what they are for in their language. Wittgenstein cited Goethe’s ‘Im Anfang war die Tat’, which he also rendered:

Our talk gets its meaning from the rest of our proceedings. (1969, §229)

<sup>5</sup> In a syntactically simple enough case, that is, and, perhaps, *ceteris paribus*.

The proceedings in which talk is embedded may make words for more than just what they are for, e.g., in English.

Within some of our activities, an object which counts as F—a leaf, say, which counts as green—is thereby to be treated, or thereby may be treated, in such-and-such ways. So treating it is part of the way those activities are to be carried out. The green leaves, say, go in that pile, or in those decorations. That idea applies to words in this way. We may define particular proceedings, or procedures, by the way they are to be carried out—by rules for their execution—and, perhaps, by the result they aim at. For example, the procedure is to compare the leaf with a given colour chart; if it matches this sample, as opposed to matching any of the others, then it goes in that pile. Such stipulated procedures are what Wittgenstein called language games, which he took as objects for comparison with actual talk in the same way that a logical calculus is an object of comparison (1953, §§81, 130–1). One comparison we might make is this. Actual words—our comments, pleas, claims, and so on—function as if they were part of certain such procedures, but not as if they were part of others: they are rightly held responsible for the role they would have in some such procedures, but not for that they would have in others; one ought to expect to be able to act on them as one would, or do with them what one would, in some, and only some such procedures. They are subject to the standards of correctness such words would have in such procedures. If words purport to say how things are, then if things are as they are according to those words, one may rely on that having the consequences for action (and for thought) that those words have, by definition, in certain stipulated procedures, or games, but not on the consequences they would have in others. Such facts are an essential part of the identification of what it is that *is* so according to the words.

To take particular words, spoken on an occasion, to be responsible for the roles they would have in certain language games is just to misunderstand them, as is failing badly enough to appreciate the responsibilities they thus have. One might, e.g., take given words ‘The leaf is green’ to have consequences for identifying its species—as not, e.g., a Japanese maple—where the words are only rightly held responsible for the role they would have in certain decorating activities. That would be a misunderstanding. Understanding what is said requires an adequate appreciation of what roles and responsibilities the words in fact take on. If given words ‘The leaf is green’ take on the roles they would in certain procedures in which such words might figure, meaning what they do, but not those they would in other such, that is an effect of their surroundings. So grasping what it is that is said in given words requires, over and above knowing what they mean, and what they refer to, an adequate appreciation of the effects of this sort which their surroundings had. I will call this a suitable sensitivity to surroundings. That now emerges as an extra capacity, beyond knowledge of meaning and referents—beyond, e.g., knowing what green is, and which leaf is in question—on which understanding what is said inescapably depends.



A capacity which consists in a suitable sensitivity to the surroundings of words is both object- and surroundings-specific. One may have it for a given remark about a leaf without being able to grasp the purposes for which given people speak of uncles as green or not, or the ways in which, within such talk, an uncle's greenness or lack of it is decided. Similarly, one may see what it would be for things to be as they are according to some remark to the effect that a leaf is green without, for the reasons just mentioned, being able to see when things would be as they are according to another such remark. Again, one may know how, within given surroundings, a leaf is to be counted as green or not, without grasping what it would be to 'go on in the same way', in those surroundings, in so classifying uncles. I do not mean that such lapses are common in daily life—though they are not unknown. But they are, in principle, possible. These specificities in the capacities involved here would make it wrong to insist on a counterpart to the generality constraint as a constraint on understanding what is said.

To sum up, grasping which object 'A' speaks of, and what it is that 'is green' speaks of in English is not enough for understanding everything, or even anything, that would be said in speaking those words, used so as to speak of what they do speak of in English. In addition, to grasp, for any particular such remark, what it would be for things to be as thus said to be, one needs a further capacity: a suitable sensitivity to the surroundings in which those words were spoken. Similarly for any predicate 'is F' which, in its language, speaks of being F. This further capacity may exist, and be adequate, for some surroundings, without existing, or being adequate, for others. It may exist, in given surroundings, for some objects one might there call F, without existing there for others. All this corresponds to these facts. First, one may know that words spoke of *a*, and of being F, while failing to grasp what it was that those words said, or what it would be for things to be as thus said to be. Second, one may understand some things which were said in calling *a* F without understanding other such things, even though one knows of all of them that they spoke of *a* and called *a* F. So, third, one's grasp of what *a* is, and what being F is, plus knowledge that words spoke of that, may put one in a position to grasp what it is that is so according to some words in which *a* was called F, without conferring grasp of that which other such words said to be so.

## V

I would like to end this essay here. I have shown that understanding what it is that given words express requires exercise of an ability independent of, so not entailed by, mere possession of the concept of being F, where the words speak of being F, plus acquaintance with the objects spoken of. Understanding involves an extra capacity, which I call sensitivity to surroundings (of the words' production). So, I would like to say, knowing what it is, in general, for a thing

to be F, plus the ability to think thoughts about *a*, yields the ability to grasp a given thought to the effect that *a* is F, only given a further capacity, which is independent of, and not guaranteed by, those just mentioned: sensitivity to the relevant way of counting things as F or not. So those first two abilities are not enough by themselves to guarantee grasp of a thought to the effect that *a* is F. They may be present without their possessor being able to grasp some such thought. The mere ability to grasp *some* thought of *b* that *b* is F is not enough to guarantee grasp of a given thought of *a* that *a* is F, even for someone capable of thinking thoughts about *a*. The extra capacity, beyond mere knowledge of what it is for a thing to be F, which yields grasp of that thought about *b* does not guarantee an equally efficacious extra capacity with respect to that particular thought about *a*. One may be suitably sensitive to what would, and what would not, show things to be as they are according to that first thought without being suitably sensitive to things being, or not, as they are according to the second. Rather, the required sensitivity is thought-specific. So the generality constraint does not hold. Q., I would like to say, e. d.

If I did stop there, though, someone would surely notice that, strictly speaking, I have shown a result only for words, and not, explicitly, for thoughts. And, it might be thought, there is no equivalent result about thoughts. It is easy to conceive of thoughts as what serve a particular function in choosing between different possible interpretations of words, and so a particular role in understanding, or misunderstanding. On the conception, to understand words aright is (at least for words which say something to be so) to associate the right thought with them—to take them to express that thought which they in fact expressed. It might then be thought that the sensitivity to surroundings whose role in understanding words I have pointed out is needed merely for seeing which of a relevant range of thoughts is the one the words in question expressed. Once that is established, there is no further role for sensitivity to surroundings. Grasp of a *thought*, the idea would be, depends on exercise of no such capacity. So it has not yet been shown that grasp of a thought requires possession, or exercise, of any capacity beyond the ones Evans indicates—in the case of a thought to the effect that *a* is F, a capacity to think about *a* plus a capacity which consists in knowledge of what it is for a thing to be F. If those are the only capacities involved in grasp of a thought, then the generality constraint may well hold.

So someone *might* think. But matters cannot quite rest there. For we still have the basic result that, for any *a*, and any way for a thing to be, F, there are *many* thoughts, each distinguished from the other in when things would be as they are according to it, *all* thoughts of *a*'s being F. Grasping some one thought to the effect that *a* is F requires a grasp of what it would be for things to be as they are according to that thought, or of when things would be that way, in contrast to ways things would be according to other thoughts of *a*'s being F. A grasp of what thus needs grasping is not provided merely by knowledge of what *a* is, and of what it is, in general, for a thing to be F. That knowledge fails to distinguish

what is so according to the thought in question from what is so according to other thoughts equally about *a* and about being F. What is needed for grasping this particular thought is adequate knowledge of when a thing would be as it is according to the particular ascription to it of being F which—insofar as the thought is about being F—is a part of the thought in question; one predication of being F, among many possible ones, according to each of which a different thing is so. One needs to grasp what it would be for a thing, and specifically, for *a*, to be F on a particular understanding of what being F would come to: namely, that understanding of being F on which, according to the thought in question, *a* is F. That understanding is not provided merely by knowledge of what being F is, since what being F is, is equally compatible with what is so according to each of the various distinct thoughts to the effect that *a* is F. Nor is the ability to grasp a particular way of distinguishing between what is F and what is not—one way among many possible ones—guaranteed merely by knowledge of what being F is.

So knowledge of what being F is, and of what *a* is, must be supplemented by an additional capacity, or capacities, if one is to grasp any particular thought to the effect that *a* is F. That much is decided by the fact that there are many thoughts to the effect that *a* is F. This already establishes the point with which I would have liked to begin this section, barring an argument that no such extra capacity, with regard to any thought to the effect that *a* is F, *could* be missing, given those first two.

We may still ask *what* extra capacities are involved in grasp of a thought.<sup>6</sup> The general point is that to grasp a particular thought to the effect that *a* is F, knowledge of what it is for something to be F, if involved at all, must be supplemented by the appropriate understanding of what a thing's being F would come to—that understanding on which, according to the thought in question, *a* is F. The extra capacity, then, is that involved in grasp of the relevant such understanding. Here is what I would like to say about that capacity. A thought, insofar as it is a thought of some *a* that it is F, is irreducibly that thought to that effect which would be thought, or expressed, in such-and-such surroundings in there predicating F of *a*. The surroundings of thinking *a* to be F play an essential and ineliminable role in determining what it would be for things to be as they are according to the thought thus thought. When things would be as they are according to the thought is not determined independently of the effect particular surroundings would have on when *a* would count as F. So grasping how things are according to the thought, insofar as it depends on grasping what it is for *a* to be F, requires appreciating which surroundings are, in this way, relevant to

<sup>6</sup> When I speak here of extra capacities, I mean capacities which, added to possession of some set of relevant concepts (including concepts of relevant objects) would yield grasp of a thought. For the moment I leave it open whether, for a given thought, there is any concept such that one could not grasp the thought without possession of that concept.

the thought and what effect such surroundings have on whether and how facts of whatever sort count towards its being so, or not, that the item, *a*, is as it is according to that thought.

But there is a rival view. On it, grasping, say, a particular thought to the effect that the leaf is green requires only grasping the right further concepts—ones of an item's being in a more finely differentiated state than merely that of being green—one, say, of an item's being the way it is if, as the leaf, it is painted green. How things are according to the thought is, on this view, just how things are if the leaf satisfies the right such finer-grained concept. The thought is that the leaf counts as green on a particular understanding of what it would be for it to be green; one grasps that understanding in grasping a rule to the effect that something is to count as green on that understanding just in case thus and so.

The original point, though, was that a concept by itself does not determine which ways for things to be, so which things, count as satisfying the concept. For any concept, there are various ways of, or techniques for, applying *that* concept, each of which would sometimes be correct; for different ways of applying it, different things would count as an item's fitting the concept—e.g., as its being green. That is the general point that the case of the leaf illustrates. If this relation between concepts and their applications is inherent in what a concept is, then there is no reason to expect the multiplicity of applications of a concept to disappear merely because we turn attention to 'more refined' concepts. Result: if someone entertains a thought in terms of some 'more refined' concept, then to grasp it, he must also grasp the right way of applying that concept—one of many possible ones. But no way of applying it *is* the right one except in suitable surroundings. One thus cannot grasp how the concept applies for purposes of *that* thought without appreciating, or being sensitive to, the effects on its application that surroundings of the relevant sort would have. So identifying thoughts in terms of finer-grained concepts is no substitute for a suitable sensitivity to surroundings.

This is the point Wittgenstein is after in his discussion of rules and what they do, which, not coincidentally, occurs within the family resemblance discussion (see esp. 1953, §§84–7). His first point is that a rule, like any other item with content, admits of various interpretations. ('No interpolation between the sign and its fulfilment does away with the sign.') If the rule is: *Do such-and-such*; then there are always various possible understandings of what doing such-and-such would come to, in any particular case in which one is to do that. The fact that the rule says to do X cannot decide which understanding of doing X is the right one; on which understanding what one thus does would be following the rule. That fact is equally compatible with *all* of these various interpretations, since they are possible ones. The fact that the rule is that one is to do X determines what one is to do in a particular case only given that the right reactions to that instruction are what they are, rather than what they might have been. Correspondingly, knowing that the rule says to do X allows one to see what to do in a particular

case only given that one is prepared to react in the right way to that instruction, and to recognize correct reactions.

Wittgenstein makes these points in the context of the family resemblance discussion to support a view of concepts. The point is that all of the facts as to which concept a given concept is, and what it is a concept of being or doing, cannot by themselves determine, for every object, and every instance of judging that object to be what that concept is a concept of, when the object would, and when it would not, count as being as thus judged to be. There are facts about what it would be, in a particular instance, for some object to count as fitting the concept, which are not derivable merely from the facts about what concept it is, or what it is a concept of, but rather depend on certain reactions being the right reactions to those facts in the case at hand. If the painted leaf counts, on an occasion, as green, that follows from no facts, by themselves, as to what it is to be green. But it does follow from the way one would properly react to those facts on a particular occasion for judging the leaf green or not. There is a family of considerations, different ones of which decide, in different cases, and different surroundings, what counts as fitting a given concept. It follows that grasping what being green is can yield grasp of a thought that something is green only given adequate sensitivity to the effects of relevant occasions for applying it on the way that concept would apply.

To grasp a thought is to grasp what it would be for things to be a certain way—that way things are according to it. To grasp that by grasping what being green is, one must grasp the appropriate way of counting things as green or not; a particular way in which facts about the leaf may count for or against its being as it is according to that thought. To grasp that is to see the relevance to the leaf's so counting of its being fit, or unfit, for such-and-such roles in such-and-such proceedings (ones of decorating, tree classifying, and so on). To grasp that is to grasp how a leaf's being green or not would be decided where certain ends were operative, certain activities under way; where things were, in such respects, as they might be on some particular occasion for counting the leaf as green or not. That sensitivity to the significance of particular ends and doings is what I have called sensitivity to surroundings. A rule is no *substitute* for a suitable sensitivity. Nor is any 'more fine-grained' concept it might characterize. Such sensitivity is needed just as much for grasp of how things are according to a thought—for grasp of that particular way for things to be which the thought is about—as it is for understanding words aright.

In any event, there is simply no such thing as 'how things are according to a thought that a leaf is green'—not, that is, if the thought is meant to be one that things are that way merely in virtue of its being to that effect, and if that is meant to be a way which, at least as a rule, things either are or are not. Equally, there is no such thing as what would, and what would not, count as a leaf's being green *independent* of the effect of given surroundings on what is to be so counted. Hence grasp of what it is for something to be green, or of what it is for a thought

to be to the effect that something is, cannot by itself yield grasp of any particular thought to the effect that something is green. Thus it is, I suggest, with grasp of thoughts in general.

## VI

It is natural to conceive of a thought as something which floats free of surroundings: a thought expressible in given surroundings is re-expressible in any, or virtually any, others. On one notion of a thought, that is nearly enough so. Any such notion, though, must give the fullest weight to this idea of Frege's:

a thought can be decomposed in many ways, and thereby now this, now that appears as subject, and as predicate. The thought itself does not yet make definite what is to be regarded as subject. If one says 'the subject of this judgement', one designates something definite only if, at the same time, one indicates a particular sort of analysis . . . One must . . . never forget that various sentences can express the same thought. (1986/1892, p. 74)

I do not claim that Frege would make of this idea what I am about to; only that what I am about to is what the general re-expressibility of thoughts requires.

Expressing the same thought in different surroundings is constrained by a principle: if you called something F, what you thus said is true just in case the way that thing was *then* counted as its being F. If you call an object green, e.g., then what the object needs to do is be some way which counts, in the surroundings of your so calling it, as its being green. Now, if I call a leaf, say, green, I *may* say what is true of it if it is painted green. I can do so in surroundings in which being green, or what I spoke of in speaking of being green, was understood in such a way that that is something a leaf can do by being painted. Suppose that I have done that. Now you want to re-express the thought I did—either by saying to be so what I did, or by saying what it is that I said to be so. You may find yourself in surroundings in which a leaf's being painted green just would not count as its being green; in which if you called the leaf green, then what you thereby said would not be true merely in virtue of its being painted green. You might be, say, writing a forester's report on the severity of the current drought. If you are in such a position, then you cannot re-express the thought I expressed in calling the leaf green by calling the leaf green. If you do call the leaf green, you will express a quite different thought, according to which a quite different state of affairs obtains.

Despite the above, we want two ideas. One is that thoughts float free of surroundings, so that the thought I expressed is still expressible, even in surroundings as different as yours. The other is that thoughts which are true under different conditions are *ipso facto* different thoughts: if A expresses a thought which would be true given thus and so, and B expresses one which

would be false given that, then A and B have expressed distinct thoughts. In the face of the fact that, on different occasions, different things would count as something's being as said to be in then calling it F, those two ideas allow us but one conclusion: to express the thought I did, in your surroundings, you must, hence can, either speak of a different item and predicate being green of it—e.g., you may speak of the outermost surface of the leaf; or you must predicate something else of the leaf. For example, you might predicate being painted green of it. The property of being painted green is not the same as the property of being green, as would emerge if, thinking you had bought green jade, you discovered that it was only painted. But if the leaf is such as you would have said it to be in describing it as painted green, it is *ipso facto* such as I said it to be in calling it green. To hold on to the two ideas, we must allow that something like this is enough of an accomplishment for re-expressing the thought I expressed. (Note how, here, what might at first seem like sloppy ordinary street practice turns out to be dictated by requirements many philosophers with a disdain for the ordinary are happy to accept.)

On this conception, it is not intrinsic to a thought that it is expressed in terms of such-and-such concepts, or by ascribing such-and-such properties, or in speaking of such-and-such objects. Rather, one and the same thought may be expressed (on suitable occasions) in deploying any of various batteries of concepts; in speaking of any of various sets of objects and ways for them to be. A thought expressed on one occasion in speaking of being F may be expressed on another in speaking of being G, G distinct from F. Frege provides us with the vocabulary we need for this conception. He distinguishes between, on the one hand, *Gedanken*, or, as he also says, *Urtheilen* (judgements), and, on the other hand, *Aussagen*, or predications. It is not intrinsic to a *Gedanke* to the effect that *a* is F that its expression involves such-and-such concepts—e.g., a concept of being F. But it is intrinsic to an *Aussage* to involve particular predicates and concepts. Unlike *Gedanken*, an *Aussage* to the effect that *a* is F can only be produced in calling *a* F. I can only produce an *Aussage* to the effect that the leaf is green by mentioning the leaf and calling it *green*. In calling a leaf green, I produce an *Aussage* and express a *Gedanke*. The *Aussage*, tied to means of expression, is thereby tied to surroundings in ways the *Gedanke* is not: what it says to be so can only be said in *that* way in suitable surroundings.<sup>7</sup>

One might think of a generality constraint either for a notion of thought modelled on *Gedanke*, or for one modelled on *Aussage*. In neither case is it plausible that such a constraint holds. Evans states that thoughts, on his conception of them, are essentially structured, and seems to take that to mean

<sup>7</sup> A predicate calculus is, of course, an *Aussagenkalkul*. It is not surprising that that should be what was developed, given that a calculus is something over which one can *calculate*.

that for each thought, there is a unique structure which is *its*: the thought is, e.g., intrinsically about a certain leaf, about being green, and about that leaf's being that. That is why Evans thinks of entertaining a thought as an exercise of a particular battery of concepts, or capacities. That idea fits the model of *Aussagen*. But the generality constraint is not plausible for *Aussagen* because of the price the intrinsicness of their structures exacts. If, e.g., the *Aussage* predicates being green of something, there is no grasping what it is for things to be as they are according to that predication without an adequate sensitivity to what, in the surroundings of that predication, the relevant item's being green would come to, or why, or how, particular facts about the item would *there* count for or against its being green—what features of the surroundings could contribute to making those facts relevant to that predication or not. The constraint fares no better if we conceive of a thought as a *Gedanke*. For then, if a thought is correctly identified, on some occasion, in some surroundings, as the thought that such-and-such is green, we cannot merely on account of that conclude that every entertaining (or grasping) of that thought involves exercise of a capacity consisting in knowing what it is for a thing to be green. For that same thought is also expressible, and thinkable, in terms of concepts quite distinct from that.

Grasp of a thought is, *inter alia*, grasp of a particular way for things to be, so of when things would count as being that way, so of how the facts may be relevant, or not, to things being that way. Spell out the way in particular terms—*a*'s being F, say—and there is no possibility of having such a grasp without sensitivity to the effects of particular surroundings. Apart from surroundings in which we may take things to be F, there is no such thing as 'the way facts are relevant to a thing's being F', or as 'those facts which are so relevant'. Until that way is spelled out in particular terms, there is nothing on to which a generality constraint might fasten.

## VII

Wittgenstein requires a thought, insofar as it is structured, to be understood in terms of a system, or systems.<sup>8</sup> There are, *inter alia*, two main thoughts behind this. One is that what is said in calling *a* F is fixed in part by an answer to the question: 'As opposed to what?' To call *a* F is, *inter alia*, to characterize it as F rather than such-and-such other things it might be. Different remarks to

<sup>8</sup> In fact, Wittgenstein speaks of *Sätze*, and not of thoughts. *Sätze* are, normally, sentences. But perhaps Wittgenstein is thinking of them as *Aussagen*. In any event, Wittgenstein's claim about systems is meant to hold, not for thoughts, but for items which are intrinsically structured in one particular way.



the effect that *a* is F rule out different alternatives to its being that, and differ accordingly in what it would be for what they say to be so. The second is that the content of a remark to the effect that *a* is F depends in part on an answer to the question 'Like what?' With what other things is *a* thus said to share something in common? Again, different such remarks about *a* require different answers. Wittgenstein also insists that neither of these requirements decides how big the systems must be in any given case.

Both these intuitive ideas derive their force, such as it is, from the thought that there are a variety of distinct things to be said in saying *a* to be F. As for the first, fixing what *a* might have been if not as said to be in calling it F would not delimit the content of any given claim that *a* is F unless different such remarks rule out different ranges of things *a* might have been instead. Lack of an answer to the 'What else?' question would not leave content indeterminate if, for given *a* and F, there were only one answer any such question could have.<sup>9</sup> Similarly for the second idea. An answer to a 'Like what?' question *chooses* between contents only if there are various sets of comparisons, different ones of which would correctly explicate different remarks to the same effect, those remarks thereby differing in content. Intuitively, that is just the point of a 'Like what?' question. Is Uncle Hugo meant to be green in the way that a salamander is green, or more in the way that punk hair is green, or perhaps in something like the way mouldy cheese is green? Different comparisons correspond to different things to be said in calling Uncle Hugo green. Both times, the point is this: if there is only one logical space in which a proposition, or thought, *could* be embedded, the space would do no genuine work in identifying content; there would be no genuine choices to be made. In fact, though, the thoughts or propositions we express, and grasp, are embedded in different logical spaces on different occasions for our thinking, and admit of various such embeddings even given some set of concepts and objects they involve.

Consideration of what drives these two ideas suggests a third. If there is a variety of distinct things to be said to be so in saying *a* to be F, then perhaps, to grasp any one of these, one must see it as belonging to a range of distinct things to be said in so speaking. We can, perhaps, grasp a particular way of counting things as being F or not only if we see it as contrasting with other possible ways, which we grasp, in which that might be done.

Though my aim here is not to defend them, nothing said here against the generality constraint counts against these ideas about systems. The generality constraint demands something quite different: to grasp any particular thing which was said, in particular surroundings, in calling some item, *a*, F, one must be able to grasp what would be said, in any surroundings whatever, in calling any item whatever F, and what would be said, in any surroundings, in then calling *a* anything one could think about at all. That is both more and less than required

<sup>9</sup> Wittgenstein developed this idea further, in a somewhat new direction, in his 1969.

for a system. If I am right about Wittgenstein, this is the very opposite of his ideas about systems. In any event, when we examine the contents thoughts have, and how one such content is distinguished from others, the idea behind the generality constraint can be seen as neither well-motivated nor correct. But we may jettison that constraint without ceasing to see each thought as, in a number of ways, importantly part of a system or systems.

# 13

## A Sense of Occasion

Some stances aim at the world: if all goes well, stance and world match. I take my pen to be on the dresser; indeed, there it is. Some stances *contain* the world: for one to *have* that stance towards something is for that thing to be so (or there). Some of these *simply* contain their object: taking them is *not* aiming at their object (which then, happily, is there to hit). John Cook Wilson saw knowledge (among other things) as a stance of this last sort. John McDowell showed why knowledge needs to be like that. But it was J. L. Austin who made the idea viable. He did it by showing how a sense of occasion is required for proper ascription of epistemic notions. This is the story of what that comes to, and why it matters.

### I. Cook Wilson's Core Conception

What is it for knowledge to contain its object? For Cook Wilson, there are two strands in that idea. The first is that knowledge is irreducible. In his terms, it is an unanalysable frame of mind. In a letter to H. A. Prichard, for example, he says

Perhaps most fallacies in the theory of knowledge are reduced to the primary one of trying to *explain* the nature of knowing or apprehending. We cannot *construct* knowing—the act of apprehending—out of any elements. (Cook Wilson 1926, p. 803)

As a special case of this general rule, knowledge, he insists (p. 100), is not a species of belief, or opinion:

There is no general character or quality of which the essential natures of both knowledge and opinion are differentiations. . . . There is no . . . common mental attitude to the object about which we know or about which we have an opinion. Moreover it is vain to seek such a common quality in belief, on the ground that the man who knows that A is B and the man who has that opinion both believe that A is B. Belief is not knowledge and the man who knows does not believe at all what he knows; he knows it.

Cook Wilson insists, in fact, that when one knows, one does *not* believe. There is more in favour of that point than may at first appear. But for present purposes I will bracket it.

The main point is that believing, or thinking so, is not a highest common factor shared by knowing and *merely* thinking so. A highest common factor would

be a frame of mind, or, more generally, condition, such that, under favourable circumstances, being in that frame of mind would just *be*, or just count as, knowing—say, that the kettle is on—while under unfavourable conditions it would, anyway, at least be such-and-such other frame of mind—say, thinking that the kettle is on, or, perhaps, one or another of such-and-such other frames of mind (supposing so, guessing so, being convinced, etc.). If there were such a condition, or frame, of mind—one which itself might or might not count as knowing, depending on further factors—then knowing would not as such involve recognition that those further factors were present. That is how things are on so-called externalist accounts of knowledge. So it is part of the view here that no form of externalism is correct.

Cook Wilson opposes highest common factor, or ‘hybrid’, accounts of knowledge in passages like these:

We want to explain knowing an object and we explain it solely in terms of the object known, and that by giving the mind not the object but some idea of it which is said to be like it—an image (however the fact may be disguised). (p. 803)

It does not do to retain *simpliciter* the statement that Logic studies *thought* and that Science studies *things*, for the very vital reason that the *formula* usually conveys the implication that things are something outside thought altogether. . . . The *thought*, not being at all supposed to be the apprehension of the *thing*, but something self-contained, something entirely mental and only apprehension of the thing as being an apprehension of some sort of a replica of the thing (i.e., when you push them to say what *knowledge of* or *thought about* the thing is). (p. 813)

Here it is not just a hybrid conception of knowledge, but a hybrid conception of thought, or thoughts, that is under attack: it is not as if a thought, say, that Pia is kind, could be factored into a thought one could have anyway, without Pia, and some relation of that thought to Pia. That is a very contemporary idea. In the case of knowledge, the point can take this form: the object of knowledge—say, its being so that  $3 + 4 = 7$ —is just part (constitutive) of that frame of mind which is knowing—say, that  $3 + 4 = 7$ . For one to be in that frame of mind is, *inter alia*, for  $3 + 4$  to equal 7. Thus does knowledge contain its object.

The second strand in Cook Wilson’s conception is the idea that knowledge is unmistakable: if  $\mathcal{N}$  knows that  $p$ , then what he sees as to whether  $p$  leaves no possibility (for him) that  $p$  is not so. To see enough of how things are to qualify as knowing that  $p$ , one must see no less than  $p$  itself. He says, for example (p. 100),

In knowing, we can have nothing to do with the so-called ‘greater strength’ of the evidence on which the opinion is grounded; simply because we know that this ‘greater strength’ of evidence of A’s being B is compatible with A’s not being B after all.

The point for the moment is this: if  $\mathcal{N}$  knows that  $p$ , then he could not have the grounds he does for taking it that  $p$  while  $p$  was not so. There is simply *no* such possibility, not even a very remote or outlandish one. We can see this as

rejecting a Lockean suggestion: that knowledge may be merely ‘certainty as great as our frame can attain to, and as our condition needs’ (Locke 332, *IV*. 11 viii). Consider an arithmetical case.  $\mathcal{N}$  grasps certain facts—that an even number is a number divisible by 2, that a prime is a number divisible only by 1 and itself. There is no such thing as matters being otherwise in those respects. And those facts exclude *all* possibility that 2 is not the only even prime. There is no deceptive counterpart or ringer case where they hold, but, for all that, there is another even prime.  $\mathcal{N}$  also grasps how those facts mean this—he sees these facts for the proof they are. He sees, say, how any even number other than 2 (or 0) would be divisible by 1 and a number other than itself—namely, 2—and so would not be a prime. So it is not so that for all he sees, he could be in a ringer case, that for all he sees there remains even the slightest possibility that 2 is not the only even prime. That what he sees does not admit of ringers is part of what he sees in seeing the facts in question as the proof they are.

Now consider a perceptual case.  $\mathcal{N}$  faces a pig, and, eyes open, sees it. Suppose  $\mathcal{N}$  grasps that he is doing that—sees what he is seeing (in this respect) for what it is.  $\mathcal{N}$  thus sees there to be a pig before him. To grasp, as he does, what he thus sees is to grasp himself to see what excludes all possibility of there being no pig (for all he can see)—to grasp himself as seeing what excludes a ringer case. If  $\mathcal{N}$  can be in that position, then it is as impossible that there may fail to be a pig for all he can see as it was in the arithmetical case that 2 might fail to be the only even prime for all he can see. If  $\mathcal{N}$  can be in this position, then, in it, he knows there is a pig. If one cannot attain a position as good as this, then, on Cook Wilson’s conception, one cannot know that there is a pig before one. Knowledge that one faces a pig, if available, is as unmistakable, secure, as arithmetical knowledge, if *that* is attainable. There is a describable position which, *if* attainable, would be knowing one faces a pig. Nothing less than this would do. There is (*pace* Locke) no second-class variety of knowledge.

Similarly for other sorts of knowledge. I do not know that Brassens was French if for all I know he might have been a closet Belgian, no matter how remote that possibility may be. As we ordinarily think of things, I *may* be in a position to dismiss such a suggestion about Brassens as nonsense. What I actually see of how things are may leave *no* such possibility, full stop. If, but only if, that ordinary view is sometimes right, one may know that Brassens was French. Knowing that, if possible at all, would be taking in that fact in no lesser way than one may take in facts of arithmetic.

On Cook Wilson’s view, then, knowledge is never merely enough, though less than absolute, certainty. Suppose we think of evidence for  $p$  as something that makes  $p$  likely—perhaps, in the best case, extremely likely—but as something that falls short of proof as strict as in the arithmetical case, something that merely gives  $p$  some probability less than 1. Thus, on his view, having even the very best evidence for  $p$  will not, so far as that goes, count as knowing that  $p$ . Some may think that having good enough evidence is a paradigm of knowing. For those,

this aspect of Cook Wilson's view may seem absurd. Austin will show why it is not. That is one face of his way of making Cook Wilson's view viable.

## II. An Accretion

Knowledge incorporates its object: if I know the toast before me to be burnt, then its being burnt is part of the frame, or state, of mind I am thus in. There is nothing in that idea as such to rule out environmental knowledge. I may know that the toast is burnt if, say, I see the burnt toast, and can make out what I thus see as burnt toast. The toast's being burnt will then be part of my frame of mind in the sense so far required. But there is a further idea in Cook Wilson that does jeopardize knowledge of one's environment. It occurs in passages like these:

The consciousness that the knowing process is a knowing process must be contained within the knowing process itself. (p. 107)

[knowledge cannot be one of] two states of mind . . . the correct and the erroneous one . . . quite indistinguishable to the man himself. [For] as the man does not know in the erroneous state of mind, neither can he know in the other state. (Ibid.)

These passages, particularly the second, insist that there is no such thing as a ringer for knowing that  $p$ —a condition which, if  $\mathcal{N}$  were in it, would be indistinguishable to him from the condition in which he is in knowing that  $p$ .  $\mathcal{N}$  would not know that  $p$  if not- $p$ . So this means that  $\mathcal{N}$  knows that  $p$ —is in the right condition—only if there is no such thing as a ringer for his situation with respect to  $p$ —a situation in which not- $p$ , but which, were he in it, he could not distinguish from his actual one in knowing that  $p$ . Conversely, in the knowing state,  $\mathcal{N}$  can distinguish his condition from all conceivable states in which not- $p$ . I will call this Cook Wilson's distinguishability principle, or, for short, *the accretion*. It insists on much more than just that knowledge must be unmistakable, or that it simply contains its object. And it spoils the account.

Given the principle, what sorts of things might one know? Arithmetical facts, one might suppose. What that would take, the idea would be, is a clear focus on the concepts involved—on what one means, or understands, say, by a number's being even, and by its being prime—together with a perspicuous organization of what one thus takes in. Being even—or what you understand by this—is something such that there is no such thing as being it without being divisible by 2. And so on. The thought is that one's mastery of one's own concepts allows one to take in such facts in such a way that there is simply no such thing as one's being wrong about them. In Freud's phrase, in such matters we are masters in our own house. If that is so, I may get myself into a ringerless frame of mind with respect to there being only one even prime. For a ringer would be a condition in which I was mistaken as to what some understanding of mine ruled out. But that my understandings do or do not rule out such-and-such is

something that is always transparent to me if I but focus. If this is right, then the accretion leaves (some) mathematical knowledge intact.<sup>1</sup> Prichard suggests that with this model for knowledge in place, its scope can be extended. For example,

Consider instances: when knowing, for example, that a noise we are hearing is loud, we do or can know that we are knowing this and so cannot be mistaken, and when believing that the noise is due to a car we know or can know that we are believing and not knowing this. (1959, p. 89)

As Prichard conceives things, for the noise *we hear* to be loud just is for things to seem to us a certain way. Things *seeming* that way is, as such, without prejudice to the condition of our environment. It may unproblematically be a part of our being in some state of mind. Some such state of mind would be ringerless. Such a state might be knowledge. Matters change, though, where the candidate object of knowledge is a matter of how the environment happens to be arranged. That is why knowledge that a car backfired is problematic.

Cook Wilson and Prichard both accept that the accretion means that we know much less than one might have thought. Prichard (p. 97) says, for example,

we are forced to allow that we are certain of very much less than we should have said otherwise. Thus, we have to allow that we are not certain of the truth of any inductive generalization, e.g., that all men are mortal, or that sugar is sweet, for we are not *certain* that anything in the nature of a man requires that he shall at some time die; we are not even certain that the sun will rise tomorrow.

What leads him to this is the idea that

when we know something we either do, or by reflecting can, know that our condition is one of knowing that thing, while when we believe something, we either do or can know that our condition is one of believing, and not of knowing: so that we cannot mistake belief for knowledge or *vice versa*. (p. 88)<sup>2</sup>

This is the accretion in a new form. The emphasis on reflection highlights the problem for environmental knowledge. Reflection alone cannot, in the nature of the case, make one privy to how one's environment is anyway—how it would be, independent of one's enjoying one's present mental states. So reflection alone cannot distinguish a state which includes some given environmental condition—the toast's being burnt, say—from one which merely seemed like that, but in which, say, it was mere illusion that the toast was carbonized—or in which there was no toast at all. The accretion thus seems to rule out knowledge of such things, at first blush to rule out environmental knowledge altogether.

<sup>1</sup> Austin and Hilary Putnam should by now have disabused us of this idea of being masters in our own house. See, e.g., Austin 1979/1940.

<sup>2</sup> For Cook Wilson's rather less tidy (but perhaps more subtle) discussion of this point, see his 1926, pp. 108–13.

The accretion installs the argument from illusion. That argument, as here understood, moves from a premiss to a penultimate, and then to an ultimate, conclusion. These moves can be expressed as follows:

Premiss: there is a ringer for  $\mathcal{N}$ 's situation with respect to  $p$  (O).

Penultimate conclusion:  $\mathcal{N}$  might be (for all he can tell/see) in a ringer situation with respect to  $p$  (O).

Conclusion:  $\mathcal{N}$  does not V that  $p$  (or does not V O) (where the values of V are such things as know (that  $p$ ), see (that  $p$ ), see O (e.g., a pig)).

When I speak of the validity of the argument, I do not mean to speak of its being licensed, or not, by some particular apodictic rules of inference. Rather, I take it, for the argument to be valid is for it to have no counter-instances—cases where the premiss is true, the conclusion, or penultimate conclusion, false. If one *rejected* Cook Wilson's conception of knowledge, then one might accept the penultimate conclusion, but reject the conclusion on grounds of some supposed properties of the concepts expressed in relevant values of V—for example, on the perhaps good ground that it may be a pig one in fact sees for all that one *may* be in a ringer situation; or, in the case of knowledge, on the supposed ground that (it is part of the concept of knowledge that) one may know that  $p$ , for all that, for all one can tell, perhaps not- $p$ . Within the present framework, though, the crucial move will be from premiss to penultimate conclusion. It is the validity of that move, in the present sense of validity, that will be at stake in what follows.

The value of 'ringer' here depends on the target of the argument. If, as at present, it is knowledge that is in question, then what matters is that  $\mathcal{N}$  should be unable, if in the ringer situation, to distinguish it from his actual one (with respect to  $p$ 's obtaining). If  $p$  is that there is a pig before him, then in the ringer, perhaps, he confronts a robotic 'pig' with artificial flesh; one cannot tell it from a pig at his distance. (You would have to be much closer and sniff carefully, or cut it open.) Again, perhaps, in the ringer, it is a peccary before him; though peccaries do not look quite like pigs, they look just like them so far as  $\mathcal{N}$  can tell—for all he knows of what a pig looks like. In the case of seeing O (a *thing* that one confronts), what matters is that in the ringer things should look just as they would, or might, if O were present. So the robotic pig will still make for a ringer. But the peccary will not.

The crucial difference can be summed up this way: what one sees is a matter of what one is responsive to; what one knows is a matter of the quality of one's responses. If a ringer is to do the work required in the case of seeing things, then it should confront  $\mathcal{N}$  with just the same things to be, by sight, responsive to. That is meant to be accomplished in this way: there is no visual feature in it but missing from the case it is a ringer for, or missing from it and present in that case—a feature whose presence or absence might be detectable



by sight, and which would show the ringer up for merely that. For a ringer to do its work in the case of knowledge, there must be nothing in the case it is a ringer for to give  $\mathcal{N}$ 's responses any credentials (relevant to knowledge) that they would not also have in the ringer case. That is meant to be accomplished in this way: anything present in the actual case that might make  $\mathcal{N}$ 's response well founded is also present in the ringer. The differences here point to a fundamental difference between seeing items one confronts, and knowing, or seeing, that; and warn against over-assimilating the first thing to those others, or vice versa.

Since it is validity that matters in what follows, it does not matter much for the present just what steps fill in the space between premiss and penultimate conclusion. But a guiding idea for filler might be this: in the ringer case, one is aware of everything one is aware of as things stand, and vice versa. For if not, one could be aware that one was in a ringer case if one was: one would fail to be aware of something one should be if it were not a ringer, or, again, be aware of something one would not be if it were not. And if, as things stand, one is not in a ringer case, one could tell that too, since one would be aware of something one would not be aware of in a ringer case (or fail to be aware of something one would be). So, for example, in the case of seeing a pig, if there is in fact a pig before me, I am aware of no more than I would be if it were only a ringer-pig. The ringer-pig would not look any different. If it did, I could tell. And similarly, supposedly, for the case of knowing that. I do not endorse such filler. The search for filler is a search for diagnosis, for what makes this argument so seductive. For the present, though, that is a secondary issue.

With the argument in force, the prospects for knowledge of one's environment are bleak. One might resist its premiss in one case or another. Norman Malcolm (1963) once suggested that it is possible to work oneself into a position such that one's understanding of what one judges in judging, say, that there is a lemon on the sideboard excludes one's proving wrong no matter what happens. If one cuts a twist out of the lemon, and, biting on it, begins to blow bubbles, or taste soap, well, some lemons on some sideboards may behave like that, given what you meant by a lemon on the sideboard. Or else (one may insist) one has suddenly begun to hallucinate. Quine counts such policy as at least not irrational.<sup>3</sup> But, all else aside, Malcolm's tactic does not deliver what we want. If I judge there to be a twist in my glass, I do not want to be judging what is compatible with my blowing bubbles if I bite on it. Nor do I want to have left the realm of knowledge. Resisting the *premiss* thus appears a futile tactic.

McDowell dismisses the accretion:

The very idea of reason as having a sphere of operation within which it is capable of ensuring, without being beholden to the world, that one's postures are all right . . . has

<sup>3</sup> See the last section of his 1961/1953.

the look of a fantasy, something we spin to console ourselves for the palpable limits of our powers. (1998/1995, p. 405)

Here McDowell is the very voice of reason. The accretion is *not* something with which we might live. But how do we earn our right to dismiss it? The idea of ringers deals only in the possible. It does not involve us with such things as talking teapots, or with any form of magic. If I am now staring down at a pig at close range, I think I do need to admit that it would be possible for there to be an animated dummy pig, so realistic that I could not tell it from the real thing by looking (even if, to my knowledge, no one has ever made one). But then just why is it not possible that my situation is a ringer, for all I can see or tell? If there are such possibilities, then with what right do I deny that I might be in such a ringer situation? At this point, we turn to Austin.

### III. Austin's Revolution

Austin joins Cook Wilson in rejecting the Lockean idea that knowledge may consist in mere 'certainty as great as our frame can attain to, and as our condition needs'. He differs, though, on where more than that is available. He says, for example,

The quotation from Locke, with which most people are said [by Ayer] to agree, in fact contains a strong *suggestio falsi*. It suggests that when, for instance, I look at a chair a few yards in front of me in broad daylight, my view is that I have (*only*) as much certainty as I need and can get that there is a chair and that I see it. But in fact the plain man would regard doubt in such a case, not as far-fetched or over-refined or somehow unpractical, but as plain *nonsense*. (Austin 1962*b*, p. 10)

On its face this is a remark about what the plain man thinks; about, in some sense, what we are prepared to say. The significance of such remarks must yet emerge. But suppose, *pro tem.*, that Austin may be in a situation where it really is just *plain* nonsense to (try to) suppose that he is not seeing a chair, or that there (perhaps) is none. Then the penultimate conclusion of the argument from illusion does not hold; so there is not that barrier to Austin's counting as knowing, in Cook Wilson's demanding sense, that there is a chair before him. If, in such a situation, someone were to suggest that Austin might be mistaken, he would be talking nonsense, *not* suggesting what was so.

This same idea—whatever exactly it comes to—occurs in Austin's comment on Geoffrey Warnock:

Warnock condemns as . . . non-minimal the form of words 'I hear a sort of purring noise', on the ground that one who says this is assuming that he isn't wearing ear-plugs; it might really be a very loud noise, which just sounds purring to him, because of the ear-plugs. But one can't seriously say to someone, 'But you might be wearing ear-plugs' *whenever* he utters that form of words; he isn't necessarily *assuming* that he isn't, he may *know*

that he isn't, and the suggestion that he might be may itself be perfectly absurd. (Austin 1962*b*, p. 138)<sup>4</sup>

It surfaces again (1962*b*, p. 114) in his critique of Ayer on incorrigibility:

. . . if, when I make some statement, it is true that nothing whatever could in fact be produced as a cogent ground for retracting it, this can only be because I am in, have got myself into, the very best possible position for making that statement—I have, and am entitled to have, *complete* confidence in it when I make it. But whether this is so or not is not a matter of what *kind of sentence* I use in making my statement, but of what *the circumstances are* in which I make it. . . . if I watch for some time an animal a few feet in front of me, in a good light, if I prod it perhaps, sniff, and take note of the noises it makes, I may say, 'That's a pig'; and this too will be 'incorrigible', nothing could be produced that would show that I had made a mistake.

I may work myself into a situation with respect to the pig before me where it is not true (because just plain nonsense) to suggest that I might be (so far as I can tell) in a ringer situation. So we may agree that for me to know that *p* is for it to be, for me, unmistakably so that *p*; and, for all that, the scope of knowledge may be roughly what we always thought. I may, for all that, know things, for example, as to where there are, and are not, pigs or chairs. Here there is a hint as to what makes things this way: the circumstances in which I am to count as knowing, or not knowing, that there is a pig before me matter. But we have yet to see just how circumstances matter and why they should.

Austin tells us

The human intellect and senses are, indeed, *inherently* fallible and delusive, but not by any means *inveterately* so.

. . . It is naturally *always* possible ('humanly' possible) that I may be mistaken, or break my word, but that by itself is no bar against using the expressions 'I know' and 'I promise' as we do in fact use them. (1979/1946, p. 98)

That the human intellect is inherently fallible should not be understood as a psychological limitation. As Austin tells us (1962*b*, p. 112),

There isn't, there couldn't be, any kind of sentence which as such is incapable, once uttered, of being subsequently amended or retracted.

Where I judge, I cannot buy immunity to error simply by judging only about certain sorts of things—certainly not where I judge about an environment. There is principle behind that (1962*b*, p. 113 n.):

But to stipulate that a sense-datum just is whatever the speaker takes it to be—so that if he *says* something different it must be a different sense-datum—amounts to making non-mendacious sense-datum statements true by *fiat*; and if so, how could sense-data be, as they are also meant to be, non-linguistic entities *of* which we are aware, *to* which we

<sup>4</sup> The comment is on Warnock (1953); see pp. 161–3.

refer, that against which the factual truth of all empirical statements is ultimately to be tested?

There is a Fregean idea here (see, e.g., Frege 1918–19). I take something to be so only where there is that which I thus take to be so—something so, or not, independent of my so thinking. Where what I take to be so need not have been so (if not elsewhere), there is at least that much room for making sense of the idea of my being mistaken. Environmental judgements, in the nature of the case, always make for that much room.

It is always possible, in this sense, that I *may* be wrong: where I take *p* to be so, that fact, so far, always leaves it open that I might be wrong. But, Austin reminds us, for it to be *possible* that I may (might) be wrong is not yet for it to be *so* that I may be. That depends, he tells us, on circumstances (in ways not yet spelt out). If it does, that makes the situation this: there are, or may be, circumstances in which, though there is, recognizably, such a thing as a ringer for my situation with respect to *p*, it is not so that I might be in a ringer situation (for all I can see or tell). If that is right, then the argument from illusion is invalid. The accretion accordingly drops out. We may, with McDowell, reject as fantasy the idea of a sphere in which knowledge coexists with the accretion. And Cook Wilson's core conception may fit with our ordinary ideas of the sorts of things one might know.

All of this, though, merely sets out the barest shape of a position. Nothing has been said about what might entitle us to it. Austin's route to entitlement is through his view of thought, and of language. At the centre of that view is his view of truth. Here is one crucial strand in that view:

There are various *degrees and dimensions* of success in making statements: the statements fit the facts always more or less loosely, in different ways on different occasions for different intents and purposes. (1979/1950, p. 130)

The idea here is a radical departure from Frege.

Truth is a particular kind of correctness, settled entirely by things being as they are. Where one thinks, or speaks, truly, there is that which one thinks, or says, to be so, which, if true, is so on any thinking, or stating of it, and, if false, is so likewise. One thinks, or speaks, truly if and only if *what* one thinks, or states, is true. Frege took that idea to impose a particular form, or standard, on thought, or talk, of what was true or false—to be reflected in particular relations between truth-bearers. That structure is, for him, unfolded in what he called the 'laws of truth', or logic. Such structure makes for standards of coherence, characteristic of, for him essential to, what is true or false. But consider a judgement that the toast is blackened. That judgement (as any other) must measure up to the way the *world* is: for it to be true, the world must be suitably arranged. No mere standard of *coherence*—governing, as it does, relations between some *truth-bearers* and others—could decide when the world would so count. Something must fill in what the standard of suitability is to be. For Frege, given *what* is judged, that

standard is provided by the notion of truth itself. For Austin, one may fix what it is that is judged, and what *truth* is, so far as that goes, leaves room for any of many standards to be in force.

For Frege, the further standard enters the picture in this way: to grasp what it is one judges in judging the toast to be blackened—what it is for it to be blackened—just is to grasp what it is for that judgement to be true. And similarly for judging any object to fit any concept.<sup>5</sup> For him, a concept maps objects on to truth values. To grasp the concept is to do no less than grasping how it does this—so in the case of the concept of being blackened (if there really is such a thing), it is to grasp just how an object must be for it to fit the concept, and so how the world must be where the concept fits. So if we think of the judgement that the toast is blackened as an application of some concept to some object,<sup>6</sup> then, for Frege, there is no substantive, or even sensible, question as to what standard *that* should come up to in order to be true—what arrangements of the world would make it so. All that is built into the judgement itself. Similarly for any other judgement.<sup>7</sup> So for Frege, for any truth-bearer, there can be only one thing it would be for it to be true, only one thing that could count as its answering, in that way, to the way things are.

In the quoted passage, Austin denies precisely this Fregean idea. Suppose someone says the toast to be blackened; it is perfectly determinate that *that* is what he said. For all that, there may be a multitude of standards by which toast may be counted as blackened or not, a multitude of ways of deciding whether given toast is precisely *that* way. For all of these standards, it is whether the toast is blackened that would be what was decided. Nothing about what it is, as such, for toast to be blackened rules any of these out *tout court*. And it may well be that by some of these, it would be true that the toast is blackened, whereas by others, it would be false (or otherwise not true). Nor have I chosen an example particularly favourable to this case. The point would hold, Austin tells us, for any way for things to be that is capturable in words of a human language—for any description of things that we can arrange for words to give. If words of a human language can express, or mention, concepts, then it applies to any concept.<sup>8</sup> If we know that someone said such-and-such to be thus and so, we do, in some sense, know what he said (to be so). The point is: whether *that* is so depends not merely on the fact that it is *this* that is to be so or not, and on the way things are, but also on what one is to count as things being that way, where this last is a

<sup>5</sup> Frege might not agree that being blackened really is a proper concept. And similarly for any other example that would illustrate Austin's point. That would just shape the difference between Austin and Frege differently.

<sup>6</sup> For Frege, it is an *Aussage*, and not a judgement (*Urtheil*) that applies a concept to an object. A judgement, or thought, for him, has such a structure only relative to an analysis. See Frege 1986/1892. But the upshot stands.

<sup>7</sup> This is the main burden of Frege's claim that truth is undefinable.

<sup>8</sup> For convenience, let us leave such things as *pure* mathematics to one side here.

genuinely substantive question. Just how black does toast need to be to count as blackened? If it is covered with Marmite, is it blackened? The concepts expressed in the words used to call the toast black do not answer such questions univocally. There is not just one thing that might count as toast being blackened. And so it is in general.

But to say that a question is left open by given factors is not to say that it is left open *tout court*. Whether a given description is true (of things as they are) is not decided by its being the description it is. It may, for all that, be decided by something. Suppose that, one fine day, someone calls a certain lake blue. There is an understanding of a lake's being blue on which, under a sunny cloudless sky, this lake would be blue. On that understanding what was said is true. There is also a way of deciding the matter by drawing a bucket from the lake and looking in it. By that standard, what was said is false. Here, then, are two standards by which the truth of what was said *might* be decided. What it is for a lake to be blue does not as such rule in favour of the one and against the other. But there are occasions on which if someone called a lake blue, and one drew a bucket and confronted him with it, what one did would be, to say the least, foolish. It would be uncomprehending. Clearly the speaker was not to be taken as speaking of its being blue in *that* way. On the reasonable understanding of his words, what he said is thus true. That may be what it is for what he said to be true full stop. Here we appeal to the circumstances of his speaking for fixing adequately a standard for truth. That is, in brief, how, for Austin, circumstances matter.

Given that circumstances play this role, there are (as a rule) many—perhaps indefinitely many—different things to be said (to be so) in saying things to be some given way (where that is a way some words of a language speak of). What produces such variety is the circumstances in which things are said to be that way. These close off some otherwise possible ways of assessing whether or not things are the way in question. They may close off enough of this for what was thus said to be determinately true, or, at worst, false—for *it* to bear no two understandings on only one of which it would have the given value. This role for circumstances flows from Austin's view of truth. It is the core of his view of language and of thought.

If circumstances work like that, then what one does say (to be so), in saying things to be such-and-such, will depend on the circumstances in which one says it.<sup>9</sup> One may (depending on one's circumstances) speak truly in saying things to be that way where, again, in some circumstances one may speak falsely in saying them to be just *that* way. Among the many things to be said to be so in saying things to be that way may be both true things and false ones, where it is up to

<sup>9</sup> There is a myth in philosophy that what a speaker in fact says in speaking given words is identical with what he *meant*, or meant to say. That would not erase the main lines of Austin's picture. It is, anyway, nothing but a myth. (If what one said is, inevitably, what one meant to say, why speak of the latter at all?)

circumstances, if anything, to decide which sort of thing (whether something true or something false) one did say on a given occasion. We may thus speak of what it would be true to say, on an occasion, as to what is, and what is not—and what would, and what would not, be—some given way. And what it is true to say about that on one occasion may not be what it is true to say on some other.

Where circumstances have such work to do, that is work that some circumstances may shirk. They may not choose between standards by which what was said would be true, and ones by which it would be false. Suppose that there are various things to be said to be so in saying some lake to be blue. Suppose that some of these are true, and some false. Suppose someone calls the lake blue in circumstances which do not decide whether what was thus said (if anything) lies among the true things, or among the false ones. Then nothing true will have been said—for all that the lake was called blue. Nor, equally, will anything false have been said. For what was said (if anything) is no more governed by standards by which it answers to the way things are than by standards by which it fails to—standards of both sorts are equally eligible, on this occasion, for deciding whether truth was told in calling the lake blue. That matter, then, is not decided. Whether circumstances do thus shirk the work required depends on what work is required—on the standards by which, in fact, the lake would be blue, and those by which, in fact, it would not. What those are depends on how the lake is. This kind of failure of circumstance is a failure to which our statements, or certainly those meant to be contingent, are always liable.

So one cannot in general say such-and-such to be such-and-such, just whenever one likes, and expect to pay no higher penalty than speaking oddly—expect, anyway, to say something either true or false. For speaking oddly can be saying nothing either true or false. What we would or would not say can thus have serious significance. If one would not use a given description in given circumstances, or of given things, that may be because using it would not then be a way of saying anything either true or false. This completes the main lines of Austin's revolutionary view of thought and language.

#### IV. Epistemology Revisited

In *Sense and Sensibilia*, Austin states the new view of thought, and applies it to epistemic notions (1962*b*, pp. 110–11):

. . . if you just take a bunch of sentences . . . impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for . . . the question of truth and falsehood does not turn only on what a sentence *is*, nor yet on what it *means*, but on, speaking very broadly, the circumstances in which it is uttered. . . . for much the same reasons there could be no question of picking out from one's bunch of sentences those that are evidence for others, those that are 'testable', or those that are 'incorrigible'. What kind of sentence is uttered as providing

evidence for what depends, again, on the circumstances of particular cases; there is no kind of sentence which *as such* is evidence-providing, just as there is no kind of sentence which *as such* is surprising, or doubtful, or certain, or incorrigible, or true.

Notions such as *evidence* depend on circumstances for truth-evaluable content in just the way notions such as *being blue* do: what counts as evidence for what depends on the occasion for saying what is evidence for what. I will now consider, briefly and sketchily, some ways in which that matters to Cook Wilson's view.

One noted feature of that view is that merely having evidence for *p*, no matter how good, never counts as knowledge—not, that is, if evidence merely weighs, no matter how heavily, in favour of *p*, but is compatible with not-*p*—not if, for all that evidence, *p* still *might* not obtain. That clashes with an intuition some will have, that if one has good enough evidence for *p*, then one knows it. Where might that intuition come from? Sid, for example, comes home reeking of a strange perfume—evidence that he is seeing someone. That empty beer bottle on the kitchen counter is evidence that Zoë, who eschews beer, is also seeing someone. Perhaps these things are merely evidence. But suppose the incident is repeated—a few days later, Sid reeks again. Or the evidence is supported by more—strange hairs on Sid's coat, mysterious receipts in the pocket, more and more late nights at the office. Then, depending on circumstances, surely Zoë is entitled to claim to *know* what is going on? And surely, all that has changed is that evidence has got better?

No doubt, one may come to *know*, in such ways, that one's partner is seeing someone. But consider the tell-tale signs—that strange perfume, the receipt for lunch at a secluded spot, Sid's suddenly depressed libido. There are at least two statuses such things might have. That scent may be evidence that Sid is seeing someone. But it may also *mean* that he is. Here 'mean' is factive: if *a* means that *b*, then given *a*, *b*. If, despite *a*, *b* does not obtain, or happen, then *a* did not mean *b* after all. Suppose Zoë sees that *a*, and recognizes it as meaning *b*. (Thus, 'recognize' being factive, *a* does mean *b*.) Then Zoë has, not mere evidence that *b*, but proof in the strictest sense: it is not just so that, for all the grounds that Zoë has, perhaps not-*b*. Zoë sees, unmistakably, that *b*. She thus knows, on Cook Wilson's conception of knowledge. Evidence is beside the point. Strange scents can be acquired in many ways—crowded elevators, over-zealous department store personnel. But suppose Sid manifestly had no such opportunities. Then in his case that strange scent may mean that he is seeing someone. If Zoë is *au fait* enough with his wonts, she may recognize the scent to mean that.

Does the scent *mean* that Sid is seeing someone? Or is it merely evidence? What does that depend on? Suppose that, though Sid *might* have got the scent through seeing someone, he also might have got it in other ways—elevators, for example. Then the scent is at best evidence. (If seeing someone is a likely way for Sid to have got the scent, then it is good evidence. If it is an unlikely enough way, then perhaps the scent is no evidence at all.) But if Sid would not have



come by the scent in any other way—if no such scenario is actually a way things might be—then the scent means that he is seeing someone. It is then up to Zoë to appreciate what is there to be appreciated.

So what means what depends on what might be, or, equally, on what would be, given such-and-such circumstances. What factors form the truths as to what might be? Here I oversimplify. But we can detect some factors. First, what might be, where that bears on factive meaning, depends on how the world is arranged. If both sheep and goats bleat, then, in general, where there is bleating, there might be sheep about, but there also might be goats. So bleating does not in general mean sheep (as opposed to goats)—though there is still room for *that* bleating to mean sheep to one expert enough to detect just that quality of bleat. Bleating might still mean sheep in goatless parts: if there are no goats in Umbria, then, perhaps, it is not so that Umbrian bleating might be from goats. But there would be goats in Umbria were they recently imported. Such *may*, sometimes, for some purposes, count as a possibility. Where it does, one speaks the truth in saying that Umbrian bleating *might* come from goats, and so does not as such mean sheep. Again, Umbrian bleating might be from goats if Umbrians had, so far, successfully concealed them (perhaps to avoid the goat tax). That *may* sometimes mean that certain bleating might be from goats. It *might* be even where it is not.

What might be thus depends, for one thing, on how things are. To that extent, it is a matter for discovery. But it also depends on what is (needs) to be treated as fixed in how things are, and what is allowed to vary. *That* is a matter liable to vary from occasion to occasion for saying what might be. *Might* Umbrians have been concealing goats? Why would one ask that? How are they to be supposed to have done it? Is the idea that they might have secret caves, or ancient ways of shrinking goats? Or is it just that no non-Umbrian has ever bothered to look in an Umbrian barn? If it is intelligible *how* this should be unknown to us, and if there is genuine uncertainty, then, perhaps, this is a way things might be. If not, not. The notions in play here—intelligibility, in the sense in which it is intelligible that a friend might be in Paris, but not intelligible, cannot be sensibly supposed, that your university may have moved there; certainty, there being reason to doubt—all these are surely occasion-sensitive. (One can *make* something uncertain (to those listening) by speaking in its favour long enough.) So, too, then, is the notion of what might be.

And so factive meaning, too, is an occasion-sensitive notion. With that idea in place, the Cook Wilsonian conception no longer clashes with clear intuition. Zoë sniffs the scent; she notes the long hairs on the collar. At once, she is certain. All becomes clear. Does she know? There are occasions for asking that question on which the scent, or the hairs, or the combination, counts as *meaning* that Sid is seeing someone. On such occasions it is not (just) evidence. On some such occasions, Zoë may count as seeing these things to mean what they thus do. For her so to count is for her to count as having proof. One may then say truly that she knows the depressing truth. There may also be occasions on

which the scent, and the hairs, count as merely evidence (though perhaps rather persuasive evidence) that Sid is seeing someone. On such occasions, Zoë counts as (understandably) certain, but as not actually knowing. Mere evidence never yields knowledge. But on our new understanding of what it is for something to be evidence, that does not mean that such things as scent and hairs cannot allow Zoë to *know* that Sid is seeing someone. They sometimes count as mere evidence; for all of that, they may sometimes count as more.

Occasion-sensitivity, once in the picture for factive meaning, is automatically in the picture for knowing as well, on Cook Wilson's core conception. If those red hairs on his lapel mean that Sid is seeing someone, and if Zoë appreciates that fact, then she knows. But whether they do mean that—or, more exactly, whether it is true to say so—depends on the occasion for saying so. Where they do not mean it, and they are all the grounds Zoë has for thinking so, then, for Cook Wilson, she does not know. Once there is occasion-sensitivity in what might be, occasion-sensitivity pervades knowing. One does not know that *p* where *p* might not be so, or where it might not for all that one can see. Whereas if one's grounds for saying so leave no such possibility, there is no cause for denying that one knows. Indeterminate or null grounds might fit the bill here—e.g., my grounds for saying that Brassens was French. Someone else does not know whether *p* if, for all his grounds, or for all he can see as to what those grounds show, *p* might fail to be so. Whereas, again, where there is no such possibility (where it is determinate enough what that would come to), he may well count as knowing whether *p*. But, again, what it is true to say as to what might be, or what might be for all *N* can see, is liable to vary with the occasion for speaking to such issues. So, too, then, is what it is true to say as to what *N* knows. I may count as knowing that Napoleon won at Jena, on some occasions for saying what I know, and for all that, fail so to count on others.

The points so far flow out of Austin's view of thought, applied, as he suggests, to epistemic notions. I have not argued that Austin's view is right. Nor will I do much in that direction here. To make the points is not to go far into a complex phenomenon, knowledge, but far enough to make Cook Wilson's core conception viable. It makes it a conception of what is recognizable to us as knowledge. Plausibly, it would not be *knowledge* that was in question if it were a notion on which we could not, sometimes, count someone as knowing that he faces a pig or a chair, or that his partner has returned home. Someone may know there is a pig before him if he sees the pig, and can take in, or make out, that this is what he does—if he can have this as his reason for taking there to be a pig. He *may* do that, provided it is not so that, for all he can tell, he might be seeing something other than a pig (or not actually seeing at all). The point now is: there might sometimes count as no such way things might be (for him, as he now is), precisely because that does not require that things *always* so count, on every occasion for considering how things now stand. For the facts as to what might be are intrinsically occasion-sensitive ones. There are not, in addition to

these, further facts as to what, occasion-insensitively, (really) might be. Just this allows it sometimes to be true (and nothing less) to say that someone, as he then is, knows there is a pig before him. If Austin has the workings of language right, it will be true to say this roughly when (outside philosophy) we would want to.

Cook Wilson's core conception could be right only if his accretion is not. Austin's view (once established) earns us our right to jettison the accretion. It earns us a right to reject the argument from illusion. For, on some occasions for stating the argument (of particular cases), its premiss may be true while its penultimate conclusion is false. Sid faces a pig—in plain daylight, and he knows as well as most of us what a pig looks like. There is plainly such a thing as a ringer for his situation. There is such a thing as facing a robotic pig with artificial flesh. Might Sid, for all he can see or tell, *be* in a ringer situation? Might he be facing a robotic pig? Perhaps yes, perhaps no, depending on the occasion for discussing Sid. Where not, the argument from illusion, as then stated, is invalid. It is not generally reliable.

What about the filler? The idea was that there should be nothing in Sid's actual situation to give his response to the pig, in taking one to be before him, any credentials relevant to knowledge that it would not also have in the ringer situation. If the ringer is a ringer, then there will be nothing, save for one thing. It may be that in the actual situation, Sid can see there to be a pig before him. If he can see this, that is good reason for taking there to be one. And, of course, he can see no such thing in the situation with the robotic pig. The way in which the argument has proved invalid leaves room for us to say this. For if it does not count as so of Sid's actual situation that it might be one of his facing a robot, for all he can see or tell, then there is no bar to saying that he does see there to be a pig before him. This point about the credentials one may have in an actual situation but lack in a ringer has been pressed hard by John McDowell—correctly. The only problem for McDowell is what entitles him to it. What is now on offer is an account of that entitlement.

Austin has thus accomplished what Cook Wilson himself did not. He has held on to the idea that knowledge does not come in various grades—various degrees of authority as to whether *p*, the lower as well as the higher. He has made good on the idea that knowledge is not merely all the certainty our condition needs, by showing how knowledge conceived as Cook Wilson conceives it is a phenomenon with the scope we recognize. Without Austin, the problem of empirical knowledge may easily seem to be this: how can one know that there are pigs about, despite the ways there undeniably are (in principle) for one to be wrong about this—the undeniable possibilities of error, for all there is for one actually to take in as to how things are? There are accounts of knowledge which are responses to that question: we take in what we do, have the grounds we have, and we do so in a hospitable environment, one in which, with those grounds, we do not run much chance of being wrong. Any form of externalism, or, more generally, highest common factor theory, is such a response. But Austin shows,

rather, how the question is misbegotten. When I know (or count as knowing) that there are pigs about, that is not despite the possibilities there are for me to be wrong: there are then (or then count as) no such possibilities, *punkt*.

## V. Corollaries

There is a kind of philosophical *mauvaise foi* exhibited, for example, in Hume's inability to maintain his philosophical beliefs when he hit the streets. One recent form of it is in the idea that there are *ordinary* knowledge ascriptions which are true enough so far as they go, but there are also the sorts of ascriptions a philosopher would make, or withhold, and these are quite another matter. So Pia may report to Max that Zoë has seen those long red hairs, and now knows that Sid is seeing someone; and Pia may thus speak truly. But if a philosopher asks whether Zoë *really* knows, then the correct answer may be that she does not. The idea will then be that the philosopher holds to higher standards, and that it is these that matter for answering *philosophical* questions. That idea may also go with this one: Pia's report is true because its content is formed by contextual factors—all she is saying is (say) that there is no *presently relevant* doubt as to whether Zoë is right—whereas what the philosopher says, or asks, depends only on what knowledge itself is—what the word 'know', as such, speaks of. Austin provides things to say about these ideas.

Taking the second idea first, suppose I were to announce, out of the blue, that Chirac is just like (or just the same as) Blair. You might well ask me what I mean by that. I might reply that I mean nothing special by it, but only what those words, as such, say: I spoke of Chirac being just like Blair, and whether I spoke truly should depend precisely and only on whether he *is* just like Blair. ('You say to me: "You understand this expression, don't you? Well, then—I am using it in the sense you are familiar with"': Wittgenstein 1953, §117.) I should not expect a sympathetic reaction. It is clear here, at least, that if I am to have said anything at all to be so, there will have to be some special content attaching to my words 'just like', thanks somehow to the circumstances in which I spoke them. If 'might be' expresses the occasion-sensitive notion I have said it to do, then knowledge ascriptions fit just this pattern. There *is* nothing one would say to be so in saying  $\mathcal{N}$  to know that  $p$ , merely by virtue of 'know' meaning what it does. Rather, for one to say something to be so, one's use of 'know' must somehow acquire special content thanks to the circumstances of one's using it. So the second idea as to what the philosopher would be doing in his special circumstances (in Hume's study) is just fantasy.

If this is so, then for a knowledge ascription to *have* content that might make it true or false, circumstances must do their job: sufficient special content must be acquired. Given the philosopher's airy lack of interest in practical affairs, there is a standing suspicion that his special circumstances will be ones that

do no such work, in which case his attempted verdicts about knowledge will be neither true nor false. One need not overplay that hand, though. Suppose that instead of 'just like' we compare 'is blue'. As we have seen, this expresses an occasion-sensitive notion *par excellence*. Still, for some descriptions of some things as blue, circumstances may not need to do much work, or at least not conspicuously. Consider a new royal blue bath towel, exhibited in plain daylight. It may be difficult to devise circumstances in which it would not be true to call this blue. Perhaps even in his study a philosopher may speak truth in so speaking. What matters here is that there are not (conspicuously) two contrasting understandings of being blue, on one of which the towel is blue, on another of which it is not. It would be drawing the wrong moral from this example if one supposed it to show that it was settled in general what the philosopher, in his study, would be saying if he called something blue (or denied that it was), or that this shows that the towel is blue by a higher standard for being blue than items which are blue on some understandings, but not on others. (Nothing so far so much as settles what 'higher' is to mean here.) That settles the second point. Perhaps, in his study, the philosopher can say *some* true things to the effect that people know this, do not know that. It does not follow that in saying what he thus does, he is speaking of knowledge in the strictest possible sense, as opposed to the rest of us, who, most of the time, speak of knowledge only in some looser sense. Again, it is not yet even fixed what this is supposed to mean.

There is a sort of worry some have had about occasion-sensitivity, both in general and for knowledge in particular. It is that the various things that are to be said truly—each on some occasion—as to whether such-and-such, where occasion-sensitivity obtained, are not mutually consistent: the (supposed) occasion-sensitive facts to the effect that things are that way do not cohere with those to the effect that they are not. That, the idea is, shows itself in this: one could not, while in a position to say some one of those things truly, consistently acknowledge the truth of some others. For example, where it is true to say that Sid does not know there are sheep behind the barn, Max cannot consistently countenance the truth of any supposed fact to the effect that Sid does know. So he cannot consistently recognize the existence of occasion-sensitivity.

The worry might be elaborated in a variety of ways. The following will do. Others are mere variants. If knowledge is occasion-sensitive, there should be situations of this form. Pia, in given circumstances, speaks truth in saying Sid to know there are sheep behind the barn: on the understanding there would then be of what knowing this would be, Sid does know. Max, (at that very time) in different circumstances, similarly speaks truth in saying Sid not to know. Max does so because of what, on his occasion, counts as what might be—the truths there are *then* to tell about this, the understanding *might be* then would bear. As it may be, on that understanding, goats might have been flooding in; some may have washed up behind the barn. But if there might be goats, then Sid

does not know there are sheep. Nothing he can see distinguishes between *those* possibilities.

If Pia spoke of *knowledge*, then what makes what she said true cannot be merely some fact about (at best) second-class knowledge. For second-class knowledge would not be knowledge on *any* understanding *knowing* bears. So what makes for the truth of what Pia said cannot be that though Sid might have been wrong, that possibility, such as it is, may be *neglected* for the purposes of what she said. Max, though, tells truth in saying there might have been goats. If there might have been,<sup>10</sup> then Sid might have been wrong. So, it seems, if Pia told the truth, then there must be a possibility to be ignored for purposes of what she said. So she cannot have spoken of knowledge. Nor can Max consistently admit she might have.

This worry is rooted in not taking occasion-sensitivity completely seriously. Several factors are at work, some endemic to occasion-sensitivity, some peculiar to the case of knowledge. I begin with the endemic. This should manifest itself as well for occasion-sensitivity in being blue. As it may be, Pia speaks truth in calling a lake blue: the way it reflects the sky on this fine day ensures that. Sid speaks truth in saying the lake not to be blue: absence of dye, copper sulphate, etc., ensures that. Suppose Sid admits that Pia spoke truth *in speaking of the lake's being blue*. Then he seems committed to this: Pia said that the lake is blue; what she thus said is true. From which it follows (the idea is) that the lake is blue, to which he is thus committed. But the truth for him to tell was meant to be that the lake is not blue. Contradiction has arrived.

It has arrived because occasion-sensitivity is allowed only half-sway. The truth Sid sees Pia to have told is that the lake is blue, on a certain understanding of its being so: that imposed by her occasion. Recognizing this truth does not commit him to what he would then say in saying the lake to be blue. For in so speaking, he would be speaking of its being so on a different understanding of what that would be. Nor is he committed to the lake's being blue *tout court*. In fact, there is no fact to that effect. So if we put it to him 'that the lake is blue', nothing so far compels him to accept that. Thus far Sid is committed to no contradiction.

This points to two senses of 'say that', and of relatives using 'that'. 'Said that *p*' (say, that the roses are red) can mean: did what would then be saying that—so saying them to be red on the understanding there would then be of their being so. Or it can mean: said what one would now (in the position of the one saying 'said that') say in saying the roses to be red—said them to be red on the understanding of their being so which that position requires. It is on the first reading only that Sid, or we, need admit that Pia said that the lake is blue. But it is the second only that licenses the move from the truth of what she thus said to that which we, or

<sup>10</sup> Hint: note the 'disquotational' move from 'Max spoke truth in saying there might have been' to 'there might have been'. This will not turn out to be legitimate.

the reporter of what she said, would say in saying the lake to be blue. There is no route here to saddling anyone with contradictory commitments. If being blue is an occasion-sensitive matter, then so far there is no problem.

Now for something specific to knowledge. If I do not know whether  $p$ , then I cannot know Sid to know this. For if not- $p$ , then Sid does not know. And for all I know, perhaps not- $p$ . This does not mean that Sid does not know, or that I am committed, on pain of contradiction, to denying that he does. (If I know Sid to be reliable, and that he says that  $p$ , I may thus know at a stroke both that  $p$  and that he knows it. So if I do not know whether  $p$ , I do not know all of these further things.) Similarly, if, from my perspective, it might be that  $p$  (say, that goats have been flooding in), I cannot *know* that there is an occasion on which this does not count as a way things might be. (I can know there is one on which one would not intuitively suppose it was a way things might be. That is a different matter.) For all that, there might be such an occasion. People better informed than me (to take one example) might enjoy it. What Pia sees might make it untrue to say what she then would in saying that caprine presence is a way things might be. To know her to enjoy such a position would be to know that those worries which need contending with in my position are, in fact, unfounded; at which point they cease to be worries that need contending with. In allowing that there might be such an occasion, though, I do not contradict what it is currently true for me to say as to what might be.

Max speaks truth in saying there might be goats. The truth he speaks is that there might be, on a certain understanding of something's being what might be: what one ought to understand by this on this occasion. Again, if there is occasion-sensitivity, then there are not, in addition to such facts as to what might be when one understands *might be* in this or that way, further facts as to what might be anyway, occasion-insensitively. It is facts of the first kind, and not such supposed further facts, that bear on the truth of knowledge ascriptions, different ones on different ascriptions. Where Sid does not know, he is not to be treated as authoritative; where he does, he is. That rule applies equally in Pia's situation and in Max's. There is no difficulty in the idea that some people, engaging with the world in given ways, ought to treat Sid as an authority while others, engaged in other ways, ought not—even if the latter cannot recognize what the former ought to do.

Max ought not to treat Sid as an authority. For he ought to treat goats behind the barn as a way things might be. If he does so treat it, he will see what Pia said as indifferent to a possibility. But what Pia's statement is indifferent to is what might be *on a certain understanding of what might be*. It need not thereby be indifferent to any way things might be, on that understanding of *might be* which its occasion calls for. (Nor is that more than Max, on his occasion, can recognize consistently.) So it need not be understood as crediting Sid with any status he might enjoy despite the existence of possibilities that he is wrong. It may be crediting him with a status he can only enjoy in having proof he grasps as

proof. What may vary from one occasion to another (from Max's, say, to Pia's) is what would count as enjoying *that*.

What, if right, would demonstrate occasion-sensitivity is this. For us, both Pia's occasion and Max's may be fully in view. We can see all that would make things count one way on the one occasion, another on the other—if the relevant notions are occasion-sensitive. If there is not occasion-sensitivity, then at most one of these occasions exhibits the facts as they really are. For there are then only occasion-insensitive facts as to what (really) might be, no matter what else passes for that on one occasion or the other. So either it really might be that there are goats behind the barn, or, really, that is not a way things might be, *punkt*. So which is it? What Austin and I think is that this question has no motivated answer. Nothing in the way things are gives the one answer any better credentials than the other as an answer to the question what (really) might be. If we are right, and if the point holds, not just for goats behind the barn, but reasonably systematically, then there can be no facts about what might be (or surely not enough) if those facts are not occasion-sensitive ones. That is always the mainspring of occasion-sensitivity. I think it is easy to confirm in the case at hand.

It is incumbent on any contextualist view of knowledge to speak, manifestly and demonstrably, of *knowledge*, and nothing less. It must not merely speak of certainty good enough for one or another condition we may be in. For that reason it will not do to speak of higher and lower, or stricter and looser, standards for knowing, as if variation is merely in what we let *pass* for knowledge. It will not do to allow that there is any occasion on which, when it comes to who counts as knowing what, there are possibilities to be ignored. I have just sketched the lines on which Austin showed how this trick is to be turned. I hope also to have indicated why the trick is needed if knowledge is to fit, as it should, Cook Wilson's core conception.

## VI. Status

John McDowell has argued compellingly that (what is in fact) Cook Wilson's core conception of knowledge had better be right. He focuses on what he calls sometimes 'the highest common factor view', sometimes 'the hybrid conception'—a view that knowledge is decomposable into some condition for which we are completely responsible, and some surrounding conditions for which we are not. On the view he opposes,

knowledge is a status one possesses by virtue of an appropriate standing in the space of reasons when—this is an extra condition, not ensured by one's standing in the space of reasons—the world does one the favour of being so arranged that what one takes to be so is so. (McDowell 1998/1995, p. 400)



The upshot . . . is that knowledge of the external world cannot be completely constituted by standings in the space of reasons. The hybrid view concedes that this is partly a matter of luck. (McDowell 1998/1995, p. 405)

He opposes, equally, the idea that knowledge may be anything less than seeing, *unmistakably*, that  $p$ :

One's epistemic standing on some question cannot intelligibly be constituted, even in part, by matters blankly external to how it is with one subjectively. For how could such matters be other than beyond one's ken? And how could matters beyond one's ken make any difference to one's epistemic standing? (McDowell 1998/1982, p. 390)

When one knows that  $p$ ,

The obtaining of the fact is precisely not blankly external to [the subject's] subjectivity, as it would be if the truth about that were exhausted by the highest common factor. (McDowell 1998/1982, p. 391)

Thus on the core points McDowell and Cook Wilson are one.

McDowell sides with Austin (against Cook Wilson and Prichard) in making the scope of knowledge (roughly) what we ordinarily suppose it to be. On his view, Cook Wilson's core conception allows one to know, say, that there is a chair a few feet in front of one. He also joins Austin in replacing talk of frames of mind with talk of statuses—in his case, standings in the space of reasons. It is easy to see how whether someone is to be accorded a certain status—expert on baseball, elder statesman, athlete—may depend on the occasion for, or point in, according it. (Even I, to my surprise, count as an expert on baseball when among those raised on cricket.) So one might well see talk of status as a special invitation to Austin's view of the occasion-sensitivity of epistemic notions. But that is an invitation McDowell declines, or at least does nothing to take up. That, I suggest, engenders problems.

If one declines the invitation, then for any thinker  $\mathcal{N}$  and anything  $p$  one might sensibly be said to know, whether  $\mathcal{N}$  knows that  $p$  depends only on how  $\mathcal{N}$  is, and how things are, or how  $\mathcal{N}$ 's environment is. It depends not at all on the circumstances in which to count  $\mathcal{N}$  as knowing or not. Thus it is ever true to say that  $\mathcal{N}$  knows that  $p$  (at a given time) if and only if it is always true to say this. And if it is ever true or false to say this, then one will always say something true or false in saying it, no matter in what circumstances. Similarly for whether  $\mathcal{N}$  might be wrong as regards  $p$ —for whether there is a possibility for him to be wrong, or that he is. These two occasion-insensitivities create (at least within the context of the core conception) a tension which, so far as I can see, has no well-motivated resolution.

Passing the free-range pig farm near Leuchars, I note, with pleasure, the happy pigs—one, in particular, his snout poking through the railings. Eschewing Austin's view, either I know, *tout court*, that a pig has poked his snout through the rails, or I do not know this, *tout court* (unless there is something odd about

this case, so that in it, as opposed to most, I neither clearly know, nor clearly do not—a possibility I rule out by *fiat*). If there are cases where someone knows such things, this case must surely be one. I may report that I saw a pig poking his snout between the rails, and it would be wrong to suggest that perhaps I do not know this. At least it would sometimes be wrong to suggest that. On the counter-Austinian view, it is sometimes wrong only if it always is.

On the other hand, if I do know, then it is not so that I might, or may, be wrong. There is simply no possibility that I am wrong. For any way for one to be wrong about a thing like that, it is not true to say that I might (may) be wrong in that way. For this not to be true to say is for this to be something it would *always* be untrue to say. So we may factor out ‘true to say’: it is flatly not so that I might (may) be wrong in that way. But it is really hard to see how things could be like this. We must admit that there are ringers for my situation with respect to that pig. It is not, so far as we know, beyond the reach of some future robotics to build robot pigs indistinguishable by sight (at my distance) from pigs, and disposed to stick their snouts through railings. If I were to suppose that such had already been done clandestinely (at the local air base) and was being tested on this ‘farm’, I would not have exceeded the bounds of sense (at least in any absolute sense). I would not have strayed from the world of possibilities into a world of pure fantasy, where teapots talk, and so on. So if whatever is possible is possible full stop, on any occasion for saying so, then, it seems, we must say that it is possible for me to be facing a robot pig, even if that is an outrageous possibility, the chances are vanishingly small, and so on. Or, if this is not so, on this particular occasion of my passing that farm, it is hard to see what a properly motivated way might be of distinguishing the cases where it is so from those where it is not (again, within our counter-Austinian framework).

It is this tension, I think, that drives some to externalism. If I, or you, dismiss the possibility of robotic pigs, with regard to my experience near Leuchars, then, within the present framework, that must appear as only amounting to saying that it is not a possibility to be taken seriously—that we need not worry about such things—and not that there really is no such possibility at all. But in that case my claim to know must stand or fall despite the (perhaps negligible) possibility that I am in a situation with robotic pigs—despite the fact that, after all, though the possibility is negligible, I might be wrong. And then, given all that I can actually see to be so, I must, after all, rely on favours from the world, even if only negligible ones, for being right. That is just what externalists maintain. It is just what McDowell (rightly) says must not be so. But it remains mysterious how we can be entitled to deny it.

I think this unresolved tension shows up in McDowell himself. For example, he says this:

We cannot eliminate what the interiorized conception . . . conceives as a quite alien factor, the kindness of the world, as a contributor to our coming to occupy epistemically

satisfactory positions in the space of reason. . . . When someone enjoys such a position, that involves, if you like, a stroke of good fortune, a kindness from the world. . . . Whether we like it or not, we have to rely on favours from the world: not just that it presents us with appearances. . . . but that on occasion it actually is what it appears to be. But that the world does someone the necessary favour, on a given occasion, of being the way it appears to be is not extra to the person's standing in the space of reasons. . . . Once she has achieved such a standing, she needs no extra help from the world to count as knowing. (McDowell 1998/1995, pp. 405–6)

The favour the world must do me if, as I find myself at time  $t$ , I am to know that  $p$ , is that 'it actually is what it appears to be'—that  $p$  actually is so, rather than (for me, then, undetectably) merely appearing to be so. But, the idea is, once that favour is done, that it has been may be part of my standing in the space of reasons: part, that is, of what it is for me to know that  $p$ . In which case, my enjoying the standing means that I need no favours.

The human intellect and senses are, to be sure, inherently fallible, though hardly inveterately so. As Austin reminds us (1979/1946, p. 98), 'a "theory of knowledge" which denies this liability' ends up 'denying the existence of "knowledge"'. But though there is this inherent liability, there is no room, I suggest, for the sort of favour McDowell here envisages. I face a pig, just on the other side of the railing from me. What I see (whether I realize this or not) is a pig before me. If I take in (register, am aware of) my doing *that*, then I have that—that I see a pig before me—as my reason for taking there to be a pig before me; in which case I have proof, in the strictest sense of proof, that there is one before me. To take this in—to *register* it, and not merely, say, to surmise it—I must have a suitable ability. I must be able to tell *when* I see a pig before me. If I have such an ability, then one favour I do not need from the world is that things 'be what they appear'. My ability is precisely one to tell how (in the relevant respect) things are. If I lack such an ability, then I cannot be said to know there is a pig before me—in the crucial instance, to register, rather than surmise. In that case, favours from the world cannot help me. So there is no room for McDowell's envisaged favours. True, any such human ability is an ability at all (alternatively, one someone has) only in a hospitable environment. I know a pig when I see one; but I could not claim this were my environment in fact rife with peccaries, and were peccaries indistinguishable (to a non-expert) from pigs. So I may, if so inclined, thank the world for providing a hospitable environment. But that is far from owing it thanks for things' being, on this particular occasion, 'what they appear'.

McDowell's picture thus makes room for a possibility that should not be there. I owe thanks to the world, in the indicated way, only if there is a possibility that things are not 'as they appear'. That it makes this room is unsurprising, since it leaves out Austin's contribution: occasion-sensitivity. Without that, things may always fail to be what they appear. For they may so fail if there may be a ringer-pig (a peccary) before me. But, without Austin, there may be, full stop, if it is ever

true to say there might. And surely *sometimes* (in philosophy, for example), it would be true to say this. Nothing in nature makes it strictly impossible for a peccary ringer-pig to be before me. Sometimes *that* fact can be expressed by saying that it is possible for me to be seeing such a ringer. So, without Austin, it is always true to say this. So (with the pig plainly before me) I might be seeing a peccary, full stop. But then I cannot know myself to see a pig before me. That result makes Austin's revolution obligatory.

McDowell speaks of cases of knowledge as our 'taking in the layout of reality'; of our being directly receptive to this. Without Austin, this can seem, and has seemed to some, to be positing some sort of supernatural capacity. Whereas if we accept Austin's transformation of Cook Wilson, it is nothing of the sort. It is simply our standing towards the way things in fact are—the actual layout of reality—in some of the ordinary, familiar ways in which we sometimes do stand to that. It may be seeing the pig before you, or being familiar with where the sheep are kept, in the very mundane ways we do these things—by looking, by having spent some time (alert and awake) on the farm. Such perfectly mundane things may just *count* as taking in the layout of reality—on the one side, the layout, on the other the looking at it. But—Austin's point—they *may* so count only on condition that sometimes they would not. I look at the pig, and now, depending on the occasion, or the purposes, for saying what I did, I may sometimes count as having seen a pig to be there, and sometimes not. What varies here is not the way the world is, but something like the point of crediting me with such a thing, or refusing to. Our ordinary accomplishments, of an 'inherently fallible' sort, may be just what it is like to take in the layout of reality. There is nothing supernatural about them. They will do for taking this in, only because whether one has taken it in or not is an occasion-sensitive matter.<sup>11</sup>

<sup>11</sup> Special thanks to Mike Martin for trenchant criticism of an earlier draft.

## References

- Austin, J. L. (1962a), *How To Do Thing With Words* (Oxford: Oxford University Press).
- , (1962b), *Sense and Sensibilia* (Oxford: Oxford University Press).
- , (1979/1940), 'The Meaning of a Word'; repr. in *Philosophical Papers*, 3rd edn. (Oxford: Oxford University Press), 55–75.
- , (1979/1946), 'Other Minds'; repr. in *Philosophical Papers*, 3rd edn. (Oxford: Oxford University Press), 76–116.
- , (1979/1950), 'Truth'; repr. in *Philosophical Papers*, 3rd edn. (Oxford: Oxford University Press), 117–33.
- , (1979/1952), 'How To Talk: Some Simple Ways'; repr. in *Philosophical Papers*, 3rd edn. (Oxford: Oxford University Press) pp. 135–53.
- Cappelen, Herman and Lepore, Ernest (2005), *Insensitive Semantics* (Oxford: Basil Blackwell).
- Clarke, Thompson (1972), 'The Legacy of Scepticism', *The Journal of Philosophy*, 69/20: 754–69.
- Cook Wilson, John (1926), *Statement and Inference* (Oxford: Oxford University Press).
- Davidson, Donald (1967), 'Truth and Meaning', *Synthese*, 17: 304–23.
- , (1983), 'A Nice Derangement of Epitaphs', in R. Grandy and R. Warner (eds.), *Philosophical Grounds of Rationality* (Oxford: Oxford University Press), repr. in Davidson, *Truth, Language and History* (Oxford: Oxford University Press, 2005), pp. 89–107.
- Dummett, Michael (1959), 'Truth', repr. in Dummett (1978), 1–28.
- , (1975), 'Wang's Paradox', *Synthese*, 30: 301–24; repr. in Dummett (1978), 248–68.
- , (1978), *Truth and Other Enigmas* (London: Duckworth).
- , (1981), *The Interpretation of Frege's Philosophy* (London: Duckworth).
- Evans, Gareth (1982), *The Varieties of Reference* (Oxford: Oxford University Press).
- , (1985), 'Understanding Demonstratives', in *Collected Papers* (Oxford: Oxford University Press), 291–321.
- Frege, Gottlob (1882a), *Die Grundlagen der Arithmetik*; trans. J. L. Austin as *The Foundations of Arithmetic* (Oxford: Basil Blackwell, 1950).
- , (1882b), 'Über die wissenschaftliche Berechtigung einer Begriffsschrift', *Zeitschrift für Philosophie und philosophische Kritik*, 81; trans. as 'On the Scientific Justification of a Conceptual Notation', in T. W. Bynum (ed.), *Conceptual Notation and Related Articles* (Oxford: Oxford University Press, 1972).
- , (1891), 'Funktion und Begriff', lecture to the Jenaischen Gesellschaft für Medizin und Naturwissenschaft, 9 Jan.; repr. in G. Patzig (ed.), *Funktion, Begriff, Bedeutung* (Göttingen: Vandenhoeck & Ruprecht, 1962), pp. 18–39; trans. as 'Function and Concepts', in P. T. Geach and M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Basil Blackwell, 1960).

- , (1986/1892), 'Über Begriff und Gegenstand', *Vierteljahrsschrift für wissenschaftliche Philosophie*, 16: 192–205; repr. in G. Patzig (ed.), *Funktion, Begriff, Gegenstand* (Göttingen: Vandenhoeck & Ruprecht), pp. 66–80.
- , (1903), *Die Grundgesetze der Arithmetik*, ii; repr. (Hildesheim: Georg Olms, 1962).
- , (1979/1915), 'My Basic Logical Insights', in *Posthumous Writings*, ed. H. Hermes, F. Kambartel, and F. Kaulbach (Oxford: Basil Blackwell), pp. 251–2.
- , (1918–19), 'Der Gedanke', *Beiträge zur Philosophie des deutschen Idealismus*, 2: 58–77; trans. and ed. Peter Geach as 'The Thought', in *Logical Investigations* (Oxford: Basil Blackwell, 1977).
- Grice, H. P. (1923), *Studies in the Way of Words* (Cambridge, MA: Harvard University Press). 'Gedankengefüge', *Beiträge zur Philosophie des deutschen Idealismus*, 3: 36–51; trans. as 'Compound Thoughts', in *Collected Papers*, ed. B. F. McGuinness (Oxford: Basil Blackwell, 1984).
- , (1980/1976), 'Meaning Revisited', repr. in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press), pp. 283–303.
- , (1989), ——— (1989/1961), 'The Causal Theory of Perception', repr. in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press), pp. 224–47.
- , (1989/1967), 'Logic and Conversation', in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press), pp. 22–40.
- Horwich, Paul (1990), *Truth* (Oxford: Oxford University Press).
- Kalish, Donald (1967), 'Semantics', in *The Encyclopedia of Philosophy* (New York: Macmillan), pp. 348–58.
- Kripke, Saul (1972), 'Naming and Necessity', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (Oxford: Basil Blackwell), pp. 253–355.
- Lewis, David (1972), 'General Semantics', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (Dordrecht: Reidel), pp. 169–218.
- Locke, John (1959/1690), *An Essay Concerning Human Understanding*, ed. Alexander Campbell Fraser (New York: Dover Publications).
- McDowell, John (1980), 'Meaning, Communication and Knowledge', in Z. van Straaten (ed.), *Philosophical Subjects* (Oxford: Oxford University Press), 117–39.
- , (1998/1982), 'Criteria, Defeasibility and Knowledge', *Proceedings of the British Academy*, 68; repr. in *Meaning, Knowledge and Reality* (Cambridge, MA: Harvard University Press), pp. 369–94.
- , (1998/1995), 'Knowledge and the Internal', *Philosophy and Phenomenological Research*, 55/4; repr. in *Meaning, Knowledge and Reality* (Cambridge, MA: Harvard University Press), pp. 395–413.
- Malcolm, Norman (1963), 'Knowledge and Belief', in *Knowledge and Certainty* (Englewood Cliffs, NJ: Prentice-Hall).
- Moore, G. E. (1993/1954), 'Wittgenstein's Lectures in 1930–1933', *Mind*, 63; repr. in Moore, *Philosophical Papers* (London: George Allen & Unwin, 1959), pp. 247–318; also in J. C. Klagge and A. Nordmann (Indianapolis and Cambridge: Hackett), pp. 46–114.
- Platts, Mark (1980), *Reference, Truth and Reality* (London: Routledge & Kegan Paul).
- Prichard, H. A. (1959), 'History of the Theory of Knowledge', in *Knowledge and Perception* (Oxford: Oxford University Press), pp. 69–199.

- Putnam, Hilary (1975/1962), 'What Theories Are Not'; repr. in *Philosophical Papers*, i: *Mathematics, Matter and Method* (Cambridge: Cambridge University Press), 215–117.
- Putnam, Hilary (1975/1963), 'Brains and Behaviour'; repr. in *Philosophical Papers*, ii: *Mind, Language and Reality* (Cambridge: Cambridge University Press), pp. 325–41.
- , (1975), 'The Meaning of "Meaning"', in *Philosophical Papers*, ii: *Mind, Language and Reality* (Cambridge: Cambridge University Press), pp. 215–71.
- Quine, W. V. O. (1961/1953), 'Reference and Modality', repr. in *From a Logical Point of View* (New York: Harper & Row, 1961), pp. 139–59.
- , (1961/1953), 'Two Dogmas of Empiricism'; repr. in *From a Logical Point of View* (New York: Harper & Row, 1961), pp. 20–46.
- , (1976), 'Quantifiers and Propositional Attitudes', *Ways of Paradox*, rev. ed. (Cambridge, MA: Harvard University Press), pp. 185–96.
- Snowdon, Paul (1980), 'Perception, Vision and Causation', *Proceedings of the Aristotelian Society*, 81: 175–92.
- Sperber, Dan, and Wilson, Deirdre (1986), *Relevance* (Oxford: Basil Blackwell).
- Strawson, P. F. (1950), 'Truth', *Proceedings of the Aristotelian Society*, suppl. vol. 24; repr. in Strawson, *Logico-Linguistic Papers*, (London: Methuen, 1971), pp. 190–213.
- Travis, Charles (1975), *Saying and Understanding: A Generative Theory of Illocutions* (Oxford: Basil Blackwell).
- , (1978), 'Meaning versus Truth'. *Dialogue* (Oct.), pp. 167–86.
- , (1981a), 'How to be a Referent', in H. Parret, M. Sbisà, and J. Verschueren (eds.), *Possibilities and Limitations of Pragmatics* (Amsterdam: Benjamins).
- , (1981b), 'Reference, Speakers and Semantics', *Language and Communication*, 1.
- , (1981c), *The True and the False: The Domain of the Pragmatic* (Amsterdam: Benjamins).
- , (1989), *The Uses of Sense* (Oxford: Oxford University Press).
- , (1994), 'On Being Truth-Valued', in S. L. Tsohatzidis (ed.), *Foundations of Speech Act Theory* (London: Routledge), pp. 167–86.
- , (2000), *Unshadowed Thought* (Cambridge, MA: Harvard University Press).
- Waismann, Friedrich (1979), *Ludwig Wittgenstein and the Vienna Circle* (Oxford: Basil Blackwell).
- Warnock, G. (1953), *Berkeley* (Harmondsworth: Penguin).
- Williamson, Timothy (1994), *Vagueness* (London: Routledge).
- , (1998), 'Indefinite Extensibility', *Grazer Philosophische Studien*, 55: 1–24.
- , and Andjelković, Miroslava (2000), 'Truth, Falsity and Borderline Cases', *Philosophical Topics*, 28/1: 211–44.
- Wittgenstein, Ludwig (1953), *Philosophical Investigations* (Oxford: Basil Blackwell).
- , (1967), *Zettel* (Oxford: Basil Blackwell).
- , (1969), *On Certainty* (Oxford: Basil Blackwell).
- , (1980), *Remarks on the Philosophy of Psychology*, i (Oxford: Basil Blackwell).
- Wodehouse, P. G. (1988), 'The Aunt and the Sluggard', in *The World of Jeeves* (New York: Harper & Row).
- Wright, Crispin (1975), 'On The Coherence of Vague Predicates', *Synthese*, 30: 326–65.
- , and Hale, Bob (1997) (eds.), *A Companion to the Philosophy of Language* (Oxford: Basil Blackwell).