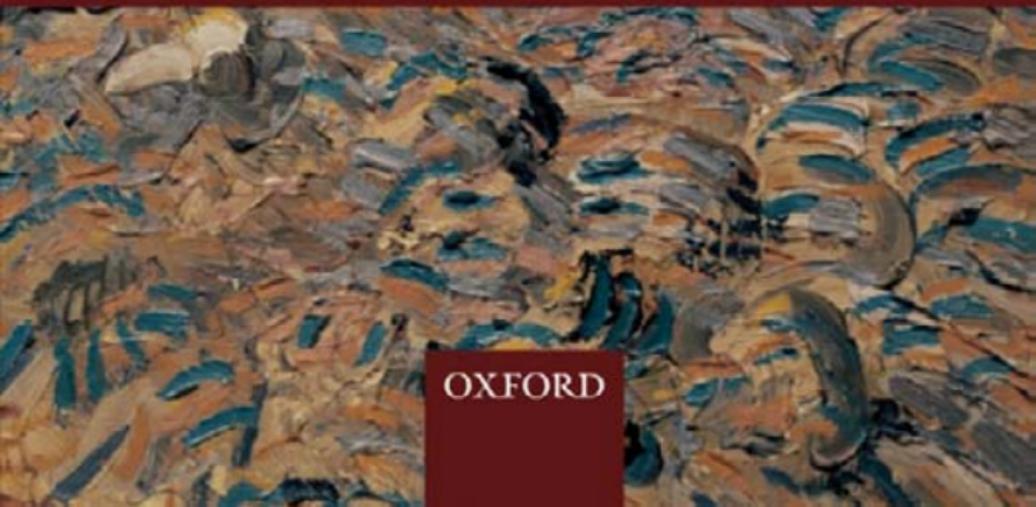




OBJECTS *of* METAPHOR

SAMUEL GUTTENPLAN



OXFORD

Objects of Metaphor

This page intentionally left blank

Objects of Metaphor

Samuel Guttenplan

CLARENDON PRESS • OXFORD

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP
Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in
Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam
Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York
© Samuel Guttenplan 2005

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)
First published 2003

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

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer
British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloging in Publication Data
Data available

ISBN 0-19-928088-6 978-0-19-928088-9
ISBN 0-19-928089-4 (Pbk.) 978-0-19-928089-6 (Pbk.)
1 3 5 7 9 10 8 6 4 2

For Jennifer

This page intentionally left blank

Contents

| | |
|---------------------------------|-----|
| Introduction | 1 |
| 1. Clearing a Space | 7 |
| 2. Object and Word | 38 |
| 3. The Semantic Descent Account | 93 |
| 4. Embellishment | 157 |
| 5. Competitors | 248 |
| Epilogue: Language and Thought | 291 |
| References | 298 |
| Index | 301 |

This page intentionally left blank

Introduction

The only reasonable response to the philosophical literature on metaphor is one of despair. This is not because what one finds there is bad; far from it. Though I don't think that anyone is quite right about metaphor, very few writers are simply wrong. Nor is the despair grounded on the sheer volume of the literature and its almost exponential increase in recent years. To be sure, if you thought you could get some grip on it in a weekend, a week, or even a month, the impossibility of this might well depress you. But of course such volume is also a sign of health in an area of investigation. It is certainly not a *reasonable* ground for desperation.

What justifiably brings one low in confronting this literature is neither its quality nor its scale, but rather the sense that, even though so many sensible things have been said about metaphor, it seems impossible to see how they might form any sort of single, coherent picture. It is as if a lot of very clever people, confronted with a huge jigsaw puzzle, all set to work in different places. Pieces, often many, are fitted together, and if you watch them being assembled, it is easy enough to share the satisfaction that comes from each additional piece snapping into place. But if you stand back to try to get some sense of the whole, what you see are only small sections, jagged in outline, which do not suggest that they themselves fit together.

What makes all the difference to the assembly of a jigsaw puzzle is what, in any actual case, is provided on the box the pieces came in: a drawing of the finished picture. In life, and here this means in our dealing with the philosophical problem of metaphor, there is no such picture. So, what we have to do is somehow fit pieces together, without forcing those that form only an 'almost' perfect fit, while at the same time sketching for ourselves at least an outline of the whole. In some areas of philosophy, one has the feeling (perhaps wrongly) that the sketch has been done, even if the pieces haven't yet been assembled; in others we have only part of that sketch. But, as the opening remark suggests, my sense is that we are in worse shape with respect to metaphor: we have many assembled sections and no real idea of how the whole ought to look. As you might expect from the use of the jigsaw metaphor, it is my aim in this book to provide the missing picture, but as you must also realize I do not think this is straightforward.

Given that philosophy is not a jigsaw puzzle in which we start with a drawing on the box, it would be wrong to think that one could provide a philosophical sketch of what needs to be done independently of trying to do it. So what follows will not be an attempt to stand to one side while offering advice about how philosophical sections of the metaphor puzzle should be joined up. That just wouldn't work. Rather, I shall offer a philosophical account of metaphor which,

though in many ways radically different from others, can highlight and absorb their best features.

This is of course an ambitious undertaking, and I admit to having had many doubts about the likelihood of its success. But I feel that I have had no choice but to persevere. Believing as I think anyone should that what is missing is a perspective on the whole—the drawing on the box—and that this can only come by providing an account into which other accounts, or at least their best features, can be fitted, a certain degree of risk-taking seemed in order. Risks here are of two kinds. On the one hand, convinced though I am that my rather different account can draw the others into a single coherent picture, its very differentness might be too radical. It might well lead to head-scratching puzzlement rather than to assent. And, on the other hand, in trying to fulfil the broader aim of gaining perspective on discussions of metaphor via my own account—going, as it were, for the bird's-eye-view—there are bound to be ways in which the account I offer might be thought under-specified, and thereby under-defended.

My strategy for coping with the first of these risks is a mixture of gradualism and minimalism. Chapter 1 sets out to show that there is indeed room for a new account of metaphor, and Chapter 2 introduces a central idea used in that account—an idea teased out of certain once-debated and now generally dormant worries about reference and, more particularly, predication. The phenomenon of metaphor will not itself figure in this discussion, the aim of which will be a certain liberation and enlargement of the role we typically associate with predicates. But my conception of predication will be crucial to everything I come to say about metaphor, so, while at first this chapter might seem a detour into philosophical logic, it will prove anything but that.¹

While the foundational work of Chapters 1 and 2 displays a necessarily defensive gradualism, the exposition of my own account of metaphor in Chapter 3 will be starkly minimalist. Using only a limited number and range of examples, the aim is to present and defend my account of metaphor with the fewest possible distractions. In this way, aside from reducing the risk of being misunderstood, it will be easiest to display its connections to other accounts. In particular, the minimalist presentation of my own account will show how it manages to assimilate many merits of these other accounts, while at the same time avoiding their defects. And this in turn will justify my claim that my account can serve as the picture on the metaphor puzzle box.

Minimalism seems the best way to display the potential of my account for giving us needed perspective, but it courts the second sort of risk mentioned above, the risk that the account, developed only with a few simple examples, will be thought unable to handle a realistically wider range of metaphors and metaphor phenomena. In the long Chapter 4, I address this worry head-on, by showing how the unadorned account of Chapter 3 accommodates complexity and richness in metaphor, and by describing how the account manages to encompass a wide variety of other

¹ In an effort to keep the reader onside, I do offer a sketch of at least the main outlines of my account at the end of Ch. 1.

phenomena of metaphor. Under this second head, I have in mind well-discussed phenomena such as dead metaphor, idiom, and simile, as well as certain less discussed ones including mixed metaphor, visual metaphor, and the figurative generally. My hope is that when the materials of Chapter 4 are added to the minimalist account in Chapter 3, you will appreciate how my account manages to organize the philosophical literature on metaphor, while at the same time telling a satisfyingly detailed story about the phenomenon of metaphor itself.² Finally, in Chapter 5, I shall consolidate these different aims by explicitly considering the relation of my account to some relatively recent philosophical work on metaphor.

So far I have been careful to say that my interest is in *philosophical* accounts of the phenomenon of metaphor, though I have said nothing about what such an account is like, or why it is the best place to start. Nor have I said anything about what counts as the phenomenon of metaphor. Both the things that I have said, and those that I haven't, require further comment.

Not that much needs to be said about the *reason* for choosing to start with a philosophical account of metaphor. Philosophers who write about metaphor might not have the subtlety of those who approach the subject from a literary point of view, nor offer the promise that psychology does of explaining metaphor's entanglement in our cognitive economy, but these latter studies depend upon, though do not provide, an account of how metaphor fits into the study of linguistic meaning generally. Yet whatever else it is, metaphor is a phenomenon of meaning, and philosophy of language has as its central aim describing, clarifying, and perhaps explaining this notion. Since I think that unless we understand metaphor's relation to meaning in general, we have not yet understood it, there can be no doubt about my choice of starting point. There is, though, a complication.

Given the centrality of meaning to the philosophy of language, and the fact that metaphor is a phenomenon of meaning, you would expect metaphor to be just as much a concern to meaning theorists as meaning is a concern to theorists of metaphor. However, reasonable though it seems, this expectation is not borne out. Though things seem to be changing, it is not unusual to find anthologies and introductory monographs in the philosophy of language which do not so much as mention the 'M'-word.³ This calls for an explanation. Indeed, given my insistence on philosophy as the right place to begin the study of metaphor, it is somewhat embarrassing to have to admit that philosophers concerned with meaning have

² Much of the detailed work of Ch. 4 should stand on its own: even if you don't accept my account of metaphor, what I say in this chapter should be of interest.

³ Look for 'metaphor' in the index of a number of introductory texts and you will see what I mean. Note too that any recent change is by and large due to the somewhat belated interest shown in Davidson's essay (1984*a*, published first in 1980) 'What Metaphors Mean'. Without denigrating the contribution that Davidson has made to the study of metaphor, it is unfortunate that his article has been the main, and sometimes the only, focus of interest. This is because it is all too easy to take away from that essay the idea that metaphor need not be prominent on the agenda of philosophy of language. This is of course a view, one that I shall consider in this book. But, with the imprimatur of Davidson himself, many have treated it, not as the first, but as the last word on metaphor.

generally neglected this phenomenon.⁴ Conveniently, in giving this explanation, it will become clearer what a philosophical account of metaphor involves, and I will be able to say something about the nature of the phenomenon of metaphor itself.

The place to begin is with some frank talk. The phenomenon of metaphor has mainly been a nuisance to philosophers of language. It is a thorn in the side of those whose aim is to give a smooth and comprehensive account of the notion of meaning. And this true whether the approach to meaning takes the route of truth conditions, communication-intentions, information theory, model theory, or any of the other current competing paradigms. Nobody adopting one or other of these views even considers the possibility that metaphor could be the starting point from which we reach out to the meanings of non-metaphorical words and sentences. And, though it wouldn't be denied that somewhere something ought to be said about metaphor, it tends either not to be said, or said late and only in passing.

Why is metaphor seen as a nuisance? Any systematic account of the meanings of words and constructions in natural language—and such accounts are the favoured candidates in contemporary philosophy of language—requires that there be some dependable link between word and meaning, or between grammatical construction and contribution to meaning. The problem with metaphor is that it seems to flout this requirement, and thereby threatens the systematicity of any account of meaning. For example, if we begin by assuming, say, that the word 'door' makes a reasonably uniform contribution to sentences in which it occurs, we will find this assumption upset by examples such as: 'The millennial fact for Hope is a large blank door that has slammed against her life' (Updike 2002: 6). In this recognizably metaphorical sentence, the dictionary meaning of 'door' as a 'hinged or sliding barrier of wood or other rigid material for closing the entrance to a building, room, carriage, safe, oven &c.' plays little or no part.

No part? Well, there are theorists of metaphor who insist that this cannot be right, and who maintain that the unruly behaviour of words in metaphors can be brought under control by complicating the story we tell about meaning. On views of this kind, there will be a route from something like the dictionary definition of 'door' to a more complex contribution this word makes to the meaning of sentences in which it occurs, a route that will take in the Updike sentence. And there are others who, not wanting to complicate the basic story about word and grammatical meaning, seek some alternative home for metaphor, such as speakers' meaning, where metaphor's power to seduce words from the straight and narrow will not be a problem.

So, on the one hand, metaphor makes a nuisance of itself in respect of our theorizing about meaning, and, on the other hand, there are philosophical accounts of metaphor which suggest ways to cope. I will consider the latter in more detail in Chapter 1, but, whatever we come to say about them, you would have expected these

⁴ Note that I am not here denying that there is a great deal written about metaphor by philosophers. The point is simply that metaphor tends not to have the place within the philosophy of language that one might expect it to have.

ways of coping to figure in the literature on meaning; after all, they seem to be contributions to meaning theory itself. Yet, as noted above, philosophers of language aiming to give an account of meaning pretty much ignore these efforts.

The reason for this is often only implicit, but what it comes down to is essentially a scepticism about whether the phenomenon of metaphor is a genuine kind with an underlying nature. I was once given the following friendly warning about getting too involved with metaphor: ‘Surely, metaphor is a rag-bag category that is best handled on the fly. You can’t think that there is anything interesting to say about it from the point of view of theory.’⁵ And unreasonable though I think such a view is, I can certainly understand why it is held.

Traditionally, metaphor was regarded as one of the so-called ‘tropes’—ways in which linguistic items are ‘turned’ or undergo conversion, so that they come to fulfil some role other than their original one. Thus, in saying that Achilles is a lion, the word ‘lion’ would have been regarded as having been turned from its usual role: whereas it is typically the name of the feline kind, it comes to have a different predicative function in respect of Achilles. As a trope, metaphor traditionally took its place in a list which included irony, meiosis, litotes, hyperbole, metonym, synecdoche, catachresis, parable, allegory, etc.—a list which seems dizzyingly extendable by anyone with the determination to make ever finer distinctions. More recently, and perhaps because the tradition of adding to the list of tropes has lost its charm, ‘metaphor’ (or, more commonly, the adjectival ‘metaphorical’ or adverbial ‘metaphorically’) turns up as a way of referring to tropes generally. An examiner claiming to have read an ‘uncountable’ number of identical examination answers is held to have spoken metaphorically, as if this adverb correctly covers the case that we can also describe as hyperbole. More liberal still is a current practice of using ‘metaphor’ pretty much interchangeably with ‘figurative’ and ‘non-literal’. Thus, we find comments like: ‘Some metaphors scarcely deserve the name ... If I tell you not to cross the path of someone who is prickly, steaming, or up in arms, you understand my language, for all its figurative nature, as immediately and certainly as if I had chosen more literal ways of expressing myself.’⁶ While this liberal usage of ‘metaphor’ is convenient, it certainly doesn’t instil confidence in the idea of metaphor as a unitary phenomenon with an underlying nature, a phenomenon robust enough to still be around when, as one will say, nature has been carved at the joints. So, as merely one among the many different kinds of trope, metaphor does not seem interesting enough to demand serious theoretical attention. But when thought of in the liberal way—especially when ‘metaphorical’ is used as more or less equivalent to ‘figurative’ and ‘non-literal’—it is difficult to imagine a defence against the charge of being a rag-bag category of no theoretical interest.

My account of metaphor aims to answer this charge. On the one hand, the account suggests that metaphor is in fact a theoretically unified, robust phenomenon which, though it casts light on notions like non-literality and the figurative, is not

⁵ It is unfair to attribute sentiments conveyed in emails, but I can say that this was from a prominent philosopher of language.

⁶ Blackburn 1984: 172. The context of the remark is not important, just the easy slide between ‘metaphor’, ‘figurative’ and ‘[non-]literal’.

interchangeable with them. On the other hand, while allowing that metaphor is a member, and not the head, of the trope family, the account will provide ample reason for thinking of metaphor as crucial to our understanding of meaning. In so far as I can convince you of these things, the neglect of metaphor in theorizing about meaning should come to seem negligent.

The list of things that I am promising is, I realize, formidable: a philosophical account of metaphor which manages to integrate other such accounts, or at least to absorb what is best in them; a justification of the idea that a philosophical account is the right place to begin the study of metaphor; a defence of the robustness of metaphor as a theoretical kind, crucial to the study of meaning; a way of distinguishing metaphor from other members of the trope family, as well as from the merely figurative and non-literal; and finally, an account which is powerful and flexible enough to accommodate the syntactic and semantic richness and diversity we expect of metaphor. However, long as the list is, there is one more thing which is in many ways the most important, even though it will not occupy much of the book.

When I began to write up my hunches about metaphor—hunches which I had as long ago as 1980—I had thought that metaphor would occupy only part of the book, and that there would be room left for a couple of chapters on what I saw as consequences of the account well outside the arena of metaphor. Moreover, it felt to me that the overall value of any contribution I was hoping to make depended on these consequences. Perhaps this was because I had unreflectively absorbed the prejudice mentioned earlier, which sees metaphor as at most an interesting feature of language, and not as central to meaning. But what seems more likely is that the material planned for these chapters—material on the supposed priority of thought over language, the course of human language development, and the debate about the linguistic abilities of non-human animals—simply struck me as exciting. Whatever the explanation, the development and defence of my account took up more space than I had originally thought, and the chapters on discussions of consequences that were planned have left only a shadow of themselves in the Epilogue. The reason that I mention this here, before I even get started, is partly that it might help to justify the space I devote in Chapter 2 to what I think of as predicate liberation—an idea fundamental to extending my account of metaphor to these other issues. More importantly, suggesting here the larger context from which this book developed might discourage you from approaching it with the prejudice about metaphor described above. Metaphor is of course a specialist topic within the philosophy of language, but, as I hope you will come to see, it is not only that.

1

Clearing a Space

1.1. Introduction

A central aim of this book is to introduce an account of metaphor alternative to those currently on offer. However, as already noted, the existing literature is vast, and accounts of metaphor abundant, so you may well think that a new one is the last thing we need. For this reason, a crucial first task will be to clear a space for my story about metaphor. The strategy is as follows: first I shall offer a principled, and I suggest exhaustive, classification of philosophical accounts of metaphor. It will turn out that the classification is tripartite, and I shall support this typology by showing how certain well-known accounts fit into it. Secondly, I will review three truths about metaphor which I believe to be non-negotiable. Non-negotiable here means that in order to be acceptable an account must accommodate each of them. There are of course other truths about metaphor, and many will be considered in later chapters. However, the three that I single out are pivotal because, as I shall argue in this chapter, one or the other of them is in fact negotiated away by the accounts typified in my classification. (There will be a certain lightness of touch in these arguments, since the mission is more that of surveying than mining.)

Demonstrating in this way the need for yet another account of metaphor is perhaps a good enough reason for the work of this chapter, but there is more to it than that. The account of metaphor I shall offer in the next chapters is not simply a revision of one or other existing account, and its very novelty could work against it. By exploring the failings of other accounts in some detail, I intend to create, not just a space for a new account, but a space of a particular shape. The hope is that you will be predisposed to see the virtues of my account when you notice how nicely it fits into this space.

1.2. Classification

A philosophical account of metaphor is a sort of balancing act. It is an attempt to balance some preferred theory of the meanings of words and constructions in a language—let me just call them ‘expressions’—with the demand that we explain in a plausible way what is going on when someone uses some such expression

metaphorically.⁷ This will be clearer with an example. With apologies for its familiarity, focus on the case of Romeo saying:(R) Juliet is the sun.⁸

According to certain philosophical theories of what Romeo's words mean, what he says is just false, even absurdly so, if not absurd *tout court*. But then an account of what he is doing cannot be based solely on what his words mean. It is just not reasonable to say that a metaphor is simply a special way of saying something false, much less that a speaker who goes in for this kind of falsehood exhibits a unique kind of irrationality. On other accounts, this is too quick. We cannot explain what Romeo is doing by a *superficial* theory of what his words mean, but a more *sophisticated* theory of meaning might yet provide a content for his words which genuinely explains what he is doing when he asserts (R). Remaining aloof from battles about which theory of meaning is correct, I shall say that an account of metaphor is 'Content Sufficient' if it appeals to some such sophisticated theory to provide an explanatory content for Romeo's words—a content that makes his linguistic act intelligible. If, on the other hand, it is held that the best theory of the meanings of Romeo's words does not yet make his act intelligible, then I shall classify this view as 'Content Insufficient'.

By way of a brief aside, it should be noted that remaining aloof from battles about the proper form of a semantic theory does not do full justice to my planned aloofness. A reasonable classification of accounts of metaphor would be made even more difficult if we had to settle the issue of where, if anywhere, to draw the line between semantics and pragmatics. Yet we neither need to be clear about such a line, nor need to know precisely what counts as pragmatics, to understand options for metaphor. Indeed, all the attempts that I have come across that base classification on these disciplinary distinctions end up casting shadows rather than light on the issues. Talk of a semantic account of metaphor, as opposed to a pragmatic one can only be useful if you take semantics to be more or less clearly circumscribed and not, as is the case, a matter of intense debate. Depending on the ways in which this debate goes, what begins as a semantic account could well end up pragmatic, or vice versa,

⁷ When I speak of 'metaphor', this is a convenient shorthand for 'the utterance of some expression or expressions identified as metaphorical'. Moreover, I have nothing important to say about the question of how we identify a particular occurrence of an expression as metaphorical. (I shall return to this point a few pages further on.) Finally, while recognizing the importance of the question of metaphor's place in the broader category of the non-literal, I shall not address this question until later in the book, preferring here (as is common) to let indisputable examples of metaphor pick out the subject matter. That said, I think the account I shall eventually offer gains support from the consequences it has for our understanding of the figurative and metaphor's place within it.

⁸ This example (from Shakespeare's *Romeo and Juliet*, I. ii. 3) is better than most in actually containing a clear and forceful example of a metaphor as opposed to the questionable Richard-is-a-lion type cases. But it is not a perfect example. It encourages the mistaken idea that metaphors are essentially of the subject-predicate form, even though, of course, it is actually syntactically and semantically more complicated than the usual S-is-P sort of case. My own account offers a way of naturally accommodating metaphors of an unlimited variety of syntactic forms, but in accord with the minimalist line described in the Introduction, I shall stick to subject-predicate cases until Ch. 4.

and this does nothing to help us understand metaphor. My notion of Content (the upper case marking a semi-technical term) seems a good way to keep the classification honest: Content is *whatever* a given theorist of meaning, using his or her best theory, attaches to the words used by a speaker. If the theorist takes an austere, truth-theoretic line about these words, then we cannot take ourselves to have understood Romeo—to have understood his linguistic act—simply by grasping that Content. But of course if a theorist favours some less austere theory of meaning—perhaps one ‘infected’ by what others would regard as pragmatic considerations—there might be some hope of this. In any case, what matters is whether Content does or does not suffice, and, as ‘Content’ is defined, this first classificatory question is not subject to the vicissitudes of choices in respect of semantic theory or decisions about the semantics/pragmatics distinction.

Allowing then, in the interest of a stable classification of theories of metaphor, a certain casualness about the line between semantic and pragmatic theories, there is more to say about Content *Insufficiency*. Clearly, a theorist who denies that Content alone is sufficient still owes us an explanation of what Romeo is doing. And there are two possibilities here. Either there is some story to be told about further, non-technical, ‘content’ which makes his act intelligible, or, though no such content is forthcoming, there is nonetheless some plausible description of the linguistic act which sends us away satisfied. I call the first option an ‘Alternative Message’ account, intending ‘Message’—admittedly a content of some kind—to be distinct from ‘Content’ which, as described, is stipulatively tied to some theory (or other) of the meaning of Romeo's words. (I use ‘Message’ because it would be too confusing to use both ‘Content’ and ‘content’ throughout.) Moreover, and unimaginatively, I call the second the ‘No Message’ account. This label of course doesn't really give us any clue about the form that any positive proposal would take, and this could make my claim of classificatory exhaustiveness look too easy. Still, while this last category is a kind of catch-all for those views that have escaped the first two categories, the ‘all’ here is very selective. In fact, such is the strength of opinion in favour of Content or Message in metaphor that only Davidson (and a handful of followers among contemporary writers) marches under the ‘No Message’ banner.

Figure 1.1 displays the classification which, at the level of generality intended, is I think exhaustive. (I have organized the diagram as a sort of flow-tree, based on answers that might be given to crucial questions.) Any account of metaphor is bound to be located at one of the three ‘Yes’ nodes. (That there isn't a third ‘No’ ruins the symmetry, but simply reflects the fact that real options run out just before a final ‘No’ could be given. After all, there isn't much point in an account of metaphor which finds (R) simply unintelligible.)

There are certainly important differences in the accounts located at each of the nodes—differences not touched upon in the typology—and, given the disputes about what counts as a theory of meaning, it may well be difficult to decide where a particular account fits. However, since the three non-negotiable truths about to be discussed are themselves insensitive to these differences and difficulties, these finer points won't matter much. Besides, it seems better to have a fixed typology, arguing

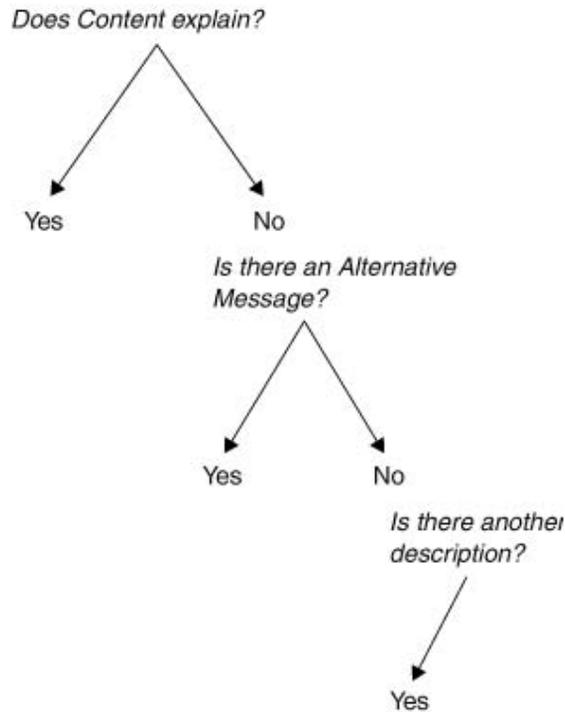


Figure 1.1

afterwards about where in it an account should be put, rather than multiplying labels in order to accommodate every different account, with the risk of losing a grip on classification.

Nonetheless, because it is important to show that the typology is useful, and that the targets of my arguments are genuine, I shall briefly consider where in the scheme to place various specific accounts of metaphor. Just before doing that, let me interpolate two brief but important notes. Both concern the innocent-sounding claim that Romeo, in uttering (R), ‘produced a metaphor’.

First, in describing (R) in this way, I might give the impression that in what follows I shall be providing an account of how speakers and hearers can *tell* whether some utterance is metaphorical. But this impression would be mistaken, for I have no theory about how we actually *identify* (R) as metaphorical, and, except for scattered anti-theoretical remarks, I shall not attempt to construct one. As is obvious enough from the literature devoted to the problem of metaphor identification, writers often take solutions to this problem to go hand-in-hand with the provision of an account of metaphor. However, I think this is unnecessary and often leads to confusion. So far from thinking that the two problems go together, it seems to me crucially important to keep them *separate*. It is one thing to give an account of metaphor—in the case at hand, to tell a story about what makes Romeo's act intelligible. It is quite another thing to give an account of how we might come to recognize, in any particular case, that we have encountered an instance of metaphor. To be sure, it can look as if the problem of accounting for metaphor is bound up with the problem of

identification because it is assumed that speakers and hearers somehow need to identify metaphors—even if not with this label—in the very act of comprehending them. Thus, the thought might be that you know you have an instance of this special phenomenon when you encounter a sincere speaker, Romeo, who says something absurdly false, namely that Juliet is the sun. And it can seem that any account you offer of this phenomenon must grow out of elements that make this identification possible. But this assumption about the relationship between identification and understanding is untenable.

On the one hand, there are ways of identifying expressions as metaphorical which have no significance for any account of the phenomenon itself. For example, sometimes we identify metaphors simply by their occurrence in a poetic context, or in the context created by a certain prose style. In these cases, there is nothing intrinsic to the expressions used that shows them to be metaphorical—it is simply that we expect metaphor in some such context. On the other hand, it may well be that a correct account of what Romeo is doing shows it to be no different from what is done with expressions we wouldn't think of as metaphorical. Or maybe, though there are differences, the expressions so understood do not serve to pick out the metaphorical, but rather some larger class which happens also to include the metaphorical. In short, the issue of metaphor identification disassociates from the problem of giving a correct account of metaphor.

It is partly because of my desire to keep the problems separate that I chose (R) as my initial example. No one doubts that Romeo's claim is metaphorical—even in the absence of any detailed account of how we come to know this. Yet this simple utterance is enough to show just how little we can be certain of, or can justifiably say, about the nature of metaphor.⁹ Of course, as the literature suggests, we are not much better off in our attempts to provide some account of metaphor identification. However, though I will not argue this at length here, this seems less a matter of philosophical difficulty than a reflection of the fact that there is no theoretical depth to the question in the first place. To repeat, my suspicion is that there are numerous, various, and not interestingly connected, ways in which we actually tell whether an utterance is metaphorical. And if this suspicion is right, then any single 'solution' to the problem of identification will seriously distort the facts. In contrast, however, I think the hunch that there is some underlying and informative story to be told about the nature of metaphor will turn out to be correct.

The second note focuses more directly on the expression 'a metaphor'. Romeo uttered the sentence (R) and it requires no particular linguistic acumen to see that in doing so he produced a metaphor. Does this mean that some of the individual words in (R) are metaphorical, that is, have special metaphorical meaning, and that is why (R) is a metaphor? Or does it mean that the effect of the whole of (R), uttered in the relevant context, is metaphorical? Or does it mean that, when the whole

⁹ Interestingly, the fact that Romeo never existed and hence never said any such thing subliminally supports the separation of issues surrounding speaker-hearer identification of metaphor from issues centered on the nature of metaphor. Perhaps this is at least a small part of (R)'s perennial appeal in writing on metaphor.

of (R) has its metaphorical effect, this induces changes in the meanings of its contained words? Dealing with these kinds of question is crucial to the whole project of giving a philosophical account of metaphor, though it is surprising how rarely they are explicitly discussed.¹⁰ It simply doesn't follow from the fact that a sentence as used produces a metaphor that its words undergo any kind of shift of meaning. Nor does it follow from the fact that none of the words in a sentence are metaphorical that the whole sentence, as used on that occasion, is not metaphorical. Being at the very beginning of my project, I shall defer my own treatment of these questions until later. But, given that I shall be liberally availing myself of the usefully non-committal expression 'a metaphor', it seemed right to acknowledge at the outset the debt that this incurs.

1.2.1. Content insufficiency

I begin with the two well-known views easiest to fit into the typology. Both Searle (1979/1993) and Davidson (1984*a*) return a clear 'No' to the first question. Neither of them regard it as plausible that the theory of expression meaning yields a suitable Content for (R). However, Searle insists that we can use principles of interpretation outside the theory of expression meaning to understand what Romeo is doing, principles which lead us to an Alternative Message. In a passage which sums up his view, he writes: In order to understand the metaphorical utterance, the hearer requires something more than his knowledge of the language ... He must have some other principles, or some other factual information, or some combination of principles and factual information that enables him to figure out that when the speaker says, 'S is P', he means 'S is R'. (Searle 1979/1993: 89)

Davidson does not follow Searle down this or any route to an Alternative Message. Instead, he invites us to understand metaphor as essentially different from paradigm cases of communication. In the latter, and details aside, a speaker produces a contentful utterance—where this content is in principle expressible in propositional form—and, if successful, either this content (my *Content*), or some message based on it (my *Alternative Message*), is transmitted to a hearer. One can say that, in favourable and appropriate circumstances, the speaker *tells* the hearer something. This certainly looks like what is going on in Romeo's case, but here Davidson insists that appearances have deceived us. We should not understand (R) as essentially Romeo's attempt to broadcast a Content or Message, so much as his putting before us a picture or image, in this case an image of Juliet as herself the sun giving off heat and light. Any picture can of course prompt many different thoughts in us, and the image of Juliet as transformed into the sun is particularly striking. But, Davidson insists, it is a mistake to think that Romeo has used his utterance to transmit these thoughts in anything like the usual communicative way. The only thing that (R) actually says—the only Content or Message in it—is patently false. (I follow Davidson in speaking of 'pictures', or 'images', but these notions should not be taken too narrowly. I can have a picture or image of a close friend that is not the kind of thing I could even in principle carry around in my wallet.)

¹⁰ One clear exception to this stricture is White 1996.

It is all too easy to lose one's grip on the divide which separates Davidson's account from others. One often hears it said that he overstates the distinction between pictures and propositions, and that his view shades into others when we view this distinction in an appropriately relaxed way.¹¹ However, I think this eclectic approach misses the radical nature of Davidson's intention. Of course, pictures in the broadest sense may function, if I can put it this way, propositionally. But we should not be misled by this. An example that Grice uses in an early stage of his celebrated attempt to define communicative meaning can be of help here.¹²

Recall that at one point Grice imagines two ways in which a person, X, might induce the belief in Y that Y's partner has been unfaithful. One way would be for X to give Y a photograph of Y's partner intimately entwined with someone else. And a second way would be for X to make a sketch of Y's partner so entwined. According to Grice, only the second of these is genuinely an act of X telling Y that Y's partner is unfaithful because only in that case does Y appropriately rely on X's communicative intentions. After all, without a proper appreciation of these intentions, the sketch might bear only a coincidental resemblance to Y's partner, and not in fact be of that person, or it might be a joke in extremely bad taste. In contrast, however, Y does not need to have access to X's intentions to interpret the photograph. In fact, it is not even crucial in this case that it was X who revealed the photograph in the first place.¹³

Whatever one comes to think about Grice's larger project—and I am certainly not suggesting that Davidson sympathizes with it—he was surely right to discern a sharp difference between an act which induces belief by relying for its interpretation on the recognition of certain communicative intentions, and one which induces belief without any such reliance. Moreover, as Grice's example shows, the distinction between pictures and propositions is irrelevant. Both the sketch and the photograph would generally be regarded as pictures.

According to Davidson, what Romeo does in uttering (R) is something like putting a photograph before us; it prompts us to have various thoughts and to form various beliefs, but none of these depend essentially on our recognizing Romeo's intentions. To be sure, Romeo's utterance is only 'like' the Gricean photograph example. Indeed, in order to be so much as aware of the metaphoricality of the utterance, Davidson suggests that we must interpret Romeo as having said something patently false, that is, as having at some level *said something*. But, to repeat, the content of this falsehood is not what makes Romeo's act intelligible.

1.2.2. Content sufficiency

Davidson's Content Insufficient account is easy to misunderstand, and that is why I spent some time on it. The problem with Content Sufficient accounts is not so much misunderstanding as variety: there are a great many of them in the literature. Each of Beardsley (1962), Black (1954–5, 1962,

¹¹ See in this regard Davies 1982/3 and Moran 1997.

¹² The example is in Grice 1957.

¹³ This is not as true as it once was. When Grice wrote, digital manipulation was not commonplace, and nowadays we might well think of photographs as little more than sketches. Still, Grice's point is clear enough.

1979/1993), Goodman (1976), Cohen (1979/1993), and, more recently, Kittay (1987) offer importantly different accounts of how to come up with expression meanings that make metaphorical utterances intelligible. However, since these differences will not matter to my later argument, I shall take the view found in the earlier work of Black as representative.¹⁴ In outline, the idea is this: the literal meanings of the expressions used in (R) do not explain what Romeo is up to. But a certain kind of further work on the meanings of these expressions can produce a suitable Content for (R). This further effort depends on both the subject and predicate in (R), and relies importantly on the ‘interaction’ between them. More specifically, there is said to be a complex system of commonplaces associated with the predicate ‘is the sun’ which are filtered by certain features of the subject—features such as, for example, Juliet's being a human being and being in a romantic relationship with Romeo—and this interactive filtering results in a new sense for the predicate. Romeo is said in this way to be using the predicate ‘is the sun’ non-standardly or non-literally, but nonetheless in a way that could be grasped by a linguistically aware audience.

A confusion about the classification of this sort of story can arise if one mistakenly identifies Content—as in *Content Sufficiency*—with a narrow conception of literal meaning, for example, a truth-conditional account. Given that such meaning on its own does not explain metaphor, it can seem as if the interactionist is in fact a Content *Insufficient* theorist. However, the views that I listed above as Content Sufficient do not identify the meaning of the words used in a metaphor with literal meaning in this sense.¹⁵ For example, Kittay (1987: 121) writes: ‘Metaphorical meaning has to do with the capabilities of language to generate meaning. When a term is used metaphorically, a “new” sense is generated. Yet we want clearly to distinguish metaphorical sense from merely another sense such as we might locate in a dictionary.’ The general line they take—and this will be crucial to my later criticism—is that an adequate theory of meaning has to recognize, and provide an account of, a kind of complexity in the theory of expression meaning. Roughly, the idea is that the meanings of everyday words, both as understood generally and on specific occasions, are a much more elaborate matter than is allowed for in typical philosophical theories of literal meaning.

The above concludes the first stage of my strategy: I have presented what I regard as a simple yet revealing classification of accounts of metaphor, and illustrated each

¹⁴ In his last published word on metaphor before his death, he says a number of things which might tempt one to think of his view as Content *Insufficient* and, therefore, closer to Searle's. The next note explains why it would be wrong to give in to this temptation.

¹⁵ Black writes (1979/1993: 28): ‘In “Metaphor” I said—scandalizing some of my subsequent critics—that the imputed interaction involves “shifts of the meaning of words belonging to the same family or system as the metaphorical expression.” I meant, of course, a shift in the speakers’ meaning—and the corresponding hearer's meaning—what both of them understand by words as used on the particular occasion.’ It is passages like this which suggest that Black may be a Content *Insufficient* theorist like Searle. For isn't Searle's main contention that speaker meaning is what we need to appeal to in metaphor? However, for all he uses the expression ‘speaker meaning’, note that Black still appeals to what is understood by the words as used. For example, it is the meaning of ‘is the sun’ which is in play, according to interactionist theory, and that theory gives the expression ‘is the sun’ a further, non-literal, meaning.

of the three resultant types with views that can be found in the literature. Moreover, I have constructed the classification so that it rides above what I regard as unenlightening disputes about the distinction between semantic and pragmatic accounts of metaphor, and literal and non-literal meaning. I now turn to the three truths I regard as non-negotiable.

1.3. Truth

The point is straightforward and has been made many times with Davidson's view as the target. It seems right that when Romeo says that Juliet is the sun, he is offering us some content, genuinely assertoric in force, something with which it makes sense to agree or disagree in the ordinary truth-directed way. It might be tactless or pointless to tell him that he is wrong, that Juliet is not the sun. But, in many cases of metaphor, the possibility of this kind of response is wholly appropriate. My suggestion—and that of many others—is that Davidson's account cannot properly take this on board. Nor can any other Content Insufficient account which tries to explain what Romeo is doing without finding in Romeo's remark a truth-evaluable content or Alternative Message.

There are of course comebacks to this. Davidson at one point suggests that we can allow for something truth evaluable by considering the beliefs prompted when we access metaphorical images. Reverting to the analogy with Grice's example, while X does not *tell* Y anything when he shows him the photograph, Y can still come away with true or false beliefs, if he has misinterpreted the photograph or if the photograph is misleading or doctored. Similarly, we may be led to believe all sorts of truths or falsehoods by some particularly arresting metaphorical image. But for those of us who think of metaphor as assertoric, this just won't do. What needs explaining about metaphorical utterances is neither that they happen to prompt true or false beliefs, nor even that their authors had hoped that some belief or other would be prompted, but rather that they enter into the give and take of declarative conversation just like any other assertion. They can be challenged, contradicted, or assented to, and their truth is thus rightly regarded as the *responsibility* of their utterers. In contrast, when X presents the photograph to Y, X does not have an assertoric commitment to the content of the photograph; it would be plainly inappropriate for Y to claim that X was mistaken, though of course he may be justified in thinking there is something morally 'off' about X's behaviour.¹⁶

Here a second defensive strategy can be called on. How important, after all, is truth to metaphor? In seriously metaphorical contexts, like that of poetry, the

¹⁶ A reader suggested that there are ways of construing Davidson's proposal which are less focused on the notion of an image, and which offer a way of dealing with my point about the give and take of conversation. Thus, think of someone telling a joke which casts aspersions on a particular person that you know. Couldn't you respond this way: 'I know what you are saying, but there is more to the person than that?' Aside from my finding this simply an odd thing to say about a joke, if the joke is intended as a model for metaphor, it certainly wouldn't be something Davidson would allow. He is adamant that the maker of a metaphor *says nothing* beyond the literal. (More on this in the next section.)

relatively banal, true–false game of assertion is left far behind. What matters is, broadly, appreciation, and this suits the image-view quite well. Taking this further, perhaps the image-theory, or any other account that doesn't bother with truth- evaluable Messages, should be seen as dealing exclusively with the exotic end of a metaphor spectrum, and therefore as not answerable to the charge of neglecting truth.

These moves remind us of important things about metaphor, but they are just not going to make the point about truth go away.¹⁷ Not all seriously metaphorical performances are found in poetry and related contexts; quite stunning metaphors figure in our everyday comments about the weather, other people, cities, and the like. And when they so figure, they are open to the challenge of getting it wrong, not because these metaphors are unappreciated or misunderstood, but because we simply don't think their authors have hold of the truth. In any case, we should be suspicious about the claim that there is something special about poetic *metaphor* as regards truth. What is special is the poetic context, not the metaphor: we do not treat wholly literal sentences in poems as straightforward assertions. Perhaps, once we identify a context as 'poetical' in a sense broad enough to include literary works that are not strictly poems, we regard the assertoric character of the discourse as having lapsed, whether in general or in some specific way.

In spite then of various evasive tactics, the central point stands: metaphorical utterances are often assertoric, and as such are liable to be judged for truth. And it is difficult to see how any Content Insufficient account of metaphor that eschews the search for some Alternative Message can accommodate this feature of metaphor. Though there are no doubt many cases in which a speaker's commitment to truth is not the main point, any account of metaphor which rules out the possibility of such commitment from the start cannot be right.

1.4. Paraphrase

Someone who asks Romeo to paraphrase his remark about Juliet is making an inappropriate or bizarre request, not merely one difficult to fulfil. This inappropriateness is the second of the truths about metaphor that any adequate theory ought to explain. Unfortunately, the notion of paraphrase which makes this claim seem so obviously true, gets mixed up with closely related notions like saying, elucidating, or explaining what a metaphor means, and when this happens, the simple truth about paraphrase can get lost. A brief consideration of an attempt to deal with paraphrase will illustrate this point.

Davidson, perhaps more than anyone else, has insisted on the inappropriateness of paraphrase, and early on in his article he offers a simple justification for it. He writes: 'Paraphrase, whether possible or not, is appropriate to what is *said*; we try,

¹⁷ I shall return to the idea of there being a spectrum of metaphor from the 'low-octane', easily expressible, to the ineffability of the poetic.

in paraphrase, to say it another way. But if I am right, a metaphor doesn't say anything beyond its literal meaning (nor does its maker say anything, in using the metaphor, beyond the literal)' (Davidson 1984*a*: 246). The trouble is that, as I argued earlier, I do not think that Davidson *is* right in claiming that a metaphor doesn't say anything beyond its literal meaning. So, while I applaud his defence of the inappropriateness of paraphrase, I cannot accept this justification of it. However, some way further on, Davidson offers a variation on this argument, one which does not set out from the assumption that a metaphor says or means nothing beyond the literal. Allowing, for the sake of argument, talk of metaphorical meaning, he notes: 'It should make us suspect a theory that it is so hard to decide even in the case of the simplest metaphors, exactly what the content is supposed to be' (Davidson 1984*a*: 262). And he goes on to suggest that, if we accept his account, we can see clearly why this should be so. He puts it this way: 'When we try to say what a metaphor 'means' we soon realize there is no end to what we want to mention. ... How many facts or propositions are conveyed by a photograph? None, an infinity, or one great unstatable fact? Bad question. A picture is not worth a thousand words, or any other number. Words are the wrong currency to exchange for a picture. (Davidson 1984*a*: 263)

Roughly, what we are offered here is an inference to the best explanation: the gap between pictures and propositions is offered as the best explanation of the inappropriateness of paraphrase.

A number of commentators have found this argument unconvincing.¹⁸ While there are important differences between pictures and propositions, it is felt that Davidson overstates the conclusion. He begins by claiming that the *difficulty* we have in saying what a metaphor means is matched by the difficulty of saying what a picture tells us. But the conclusion of the passage above actually suggests that there is a kind of *impossibility* here—that words cannot be exchanged for a picture—and, aside from rhetoric, there seems little to support this. Of course, Davidson could retreat on this point: he could say that it is not inappropriate or impossible but just terribly difficult to exchange words for pictures. And this concession would turn his overall argument into an inference supporting, not the outright inappropriateness of paraphrase, but only the difficulty of saying what a metaphor means. I think such a retreat would be a mistake. There is something right about the stark inappropriateness of paraphrase, and, separately, something right about the 'mere' difficulty we have in saying what metaphors mean. What is not so clear is how both of these can be supported by a single argument built around the supposed distinctness of pictures and propositions.

The key to getting these points correctly aligned lies in keeping separate the notion of, on the one hand, paraphrasing an utterance, and, on the other, saying what it tells us, or means. These are easily confused, not least because they sometimes come to the same thing, but they are in the end quite different notions. To take the case most closely related to Davidson's view, imagine that you are given

¹⁸ I have in mind here Davies 1982/3 and Moran 1997.

a photograph and asked: what does this tell you? This is not an easy question to answer and, as Davidson suggests, the task is endless. Still, it is certainly neither inappropriate nor impossible to make the attempt. If, in the earlier passage, Davidson is really insisting on impossibility, then his critics are surely right. Suppose instead, however, that you were asked to paraphrase the photograph? Our first reaction—I think the correct one—is that this just cannot be done. Here the trouble is not that the task is so difficult as to approach impossibility. Nor can it be blamed on the fact that photographs are somehow mute, since we have seen that it is perfectly reasonable to count them as telling us things. Rather, the bizarre nature of this request comes from our understanding paraphrase, not simply as trying to *say* the same thing in another way, but as trying to do it *in other words*.¹⁹ Since a photograph, whatever it tells us, is not itself in words, it is inappropriate to paraphrase it.

The notion of paraphrase, albeit in scare-quotes, figures explicitly the discussion leading up to Davidson's impossibility argument, but, in that argument, he implicitly takes 'paraphrasing X' to be more or less equivalent to 'saying what X tells us or means'. But, as noted, these notions are not the same; one can say what something means without being counted as paraphrasing it. Ironically, if Davidson had kept these notions distinct, his argument would have achieved both aims, without offering a target to critics. As has been noted, his imagistic or picture account of metaphor certainly makes intelligible why paraphrase, strictly understood, is out of the question, and even critics of the 'wrong currency' claim would accept that it is certainly difficult to say what a picture or photograph tells us.

I have charged Davidson with allowing paraphrase, properly understood, to drop out of his story. But, where metaphor is concerned, someone might feel that this is unfair. After all, Romeo does in fact use words, and, when we attempt to say what he means, we use words too. So why isn't the latter an attempt at paraphrase in the strictest sense? Photographs clearly are not 'paraphrased' when we try to put what they tell us into words, but this might be thought a relatively trivial fact that doesn't carry over to overtly linguistic examples.

An example will show why this is not right. Consider a code which uses what appear to be perfectly ordinary English sentences but in which the 'real' meanings are quite different, and require reference to a decoding key. For instance, suppose someone says: 'The cat is on the mat', and suppose also that the key to the code correlates this with: 'The money is in the bank'. Clearly, with access to the key, it is easy enough to say what the coded sentence means. But are we thereby *paraphrasing* the coded sentence? Surely not, and the reason is clear enough from the definition cited earlier. A paraphrase is a 're-statement of the sense of a passage in other words', and, though the coded sentence uses words in English, these words are not used to state something which is then re-stated in the decoded sentence.

Nor is this an isolated or trivial example. Indeed, it has a bearing on something which has puzzled commentators about metaphor, namely the relative ease with

¹⁹ 'Re-statement of the sense of a passage in other words' is how the *Oxford English Dictionary* puts it.

which utterances like Romeo's can be rendered into other languages.²⁰ If one is not careful to distinguish paraphrase from closely related notions, it can seem odd that, for all the trouble metaphors give us in regard to paraphrase, we often have little or no difficulty in translating them. Still, for reasons brought out by the code example, I suggest that it is simply wrong to regard translation as paraphrase. For a non-English speaker, Romeo's words do not state anything; they are just like a sentence in a code, one for which there is as yet no key. Seen in this way, a translation decodes the sentence for the non-English speaker, and it therefore does not consist in a 're-statement' of the sense of the original, so much as a first presentation of it. Thus, in at least straightforward cases of translation, it simply doesn't count as paraphrase, even though it is certainly a way of conveying to a particular audience what a sentence means.

Suppose I have convinced you that, in the strict sense of the term, it is simply inappropriate to try to paraphrase a metaphor. What does this fact do for us, and is it really so important? Isn't the really interesting issue the one of trying to say what metaphors mean? Thus, even if Davidson didn't keep his eye firmly on paraphrase, isn't the most pressing need the one for an explanation of our apparent inability to say what metaphors mean, whether we call this paraphrase or not?

One can only really decide whether a fact is 'important' against a theoretical background that puts the fact into a larger framework. As yet the only such background I have considered is Davidson's account, and I have previously found reason independent of the present issue to be dissatisfied with it. Still, the very fact that we do distinguish between paraphrasing a metaphor and saying what it means seems to me an interesting starting point. If there were to be some comprehensive account explaining the fact that we treat the first as inappropriate and second as merely difficult then what appears as a 'small' truth about the way we use 'paraphrase' will take on a much deeper significance. I intend to provide such an account in Chapter 3, but for the present it is worth observing just how much difficulty various accounts of metaphor have with this small truth.

Consider first Searle's Alternative Message account. So far from recognizing its inappropriateness, Searle's story depends for its adequacy on allowing for the possibility of paraphrase. Using one of his examples, when someone says: (1) Richard is a gorilla, and intends this metaphorically, Searle takes these words to express, not a literal Content, but some metaphorical Alternative Message. Given this, he claims it to be perfectly intelligible to be asked to say *in other words* what (1) means when understood

²⁰ I say 'relative ease' because of course there are cases where 'direct' translation of a metaphor is problematic. Pinning down the source of these problems turns in large part on getting clear about idiom, dead or frozen metaphor, and metaphor itself. Further discussion of this will figure in Ch. 3, but the point in the text is simply that there is a prima-facie mismatch between difficulty in paraphrase and translation.

metaphorically. In fact, he helpfully gives us the result of this search and calls it the 'literal paraphrase' of (1):(1) Richard is fierce, nasty and prone to violence.

Searle never suggests that paraphrase is always this easy. Noting that there are metaphors for which the search is difficult, perhaps endless, he seems to think that this concession protects him from the charge of oversimplifying metaphor. However, my complaint is not that he makes paraphrase seem easy, so much as that he makes it possible. Nor can this complaint be deflected by substituting some looser notion of saying what a metaphor means for paraphrase strictly understood. In contrast to Davidson's, there is simply no room in Searle's account for making this distinction.

Searle treats metaphor as a use of words to say something, albeit something that would be missed by anyone who took the words literally. When we attempt to capture what is said—to say in our own words what the speaker's use of words means on that occasion—we can be correctly described as re-stating what was originally expressed. Thus, unlike the cases of working out what a photograph tells us or breaking some code, attempting to say what a metaphor means is for Searle strictly a way of paraphrasing it. An accurate parallel would be a case of trying to put a piece of legal jargon into plain English: there is here no distinction between saying what it means and paraphrasing it. But this shows that Searle's account simply has no resources with which to explain why it is bizarre to be asked to paraphrase a metaphor.

Various Content Sufficient accounts fare no better on this issue. If, as Black does, you think that the words in Romeo's utterance have, in addition to literal meaning, a metaphorical one, then it is reasonable to find it difficult to say in particular cases what this meaning is. This is because metaphorical meaning is an *elaboration* of the literal and, as such, it might well be difficult to recover. However, the story about elaboration leaves us with no way to justify the blank inappropriateness of asking Romeo, or anyone else, to paraphrase his remark. We could of course just stipulate that metaphors are such a special use of language that they cannot, in the strict sense of the term, be paraphrased, even though we can make some stab at saying what they mean. But this only puts the problem back a bit, for we still need some explanation of this specialness. Whatever its defects, Davidson's account did provide this explanation: if metaphorical utterances put something like pictures in front of us, it is easy to see why we find it difficult, though not impossible, to express what they convey, and impossible, not merely difficult, to paraphrase them.

A final note: I have been careful to say that only certain Content Sufficient accounts are subject to this objection. Having at this point claimed that we must reject both No Content (e.g. Davidson) and Alternative Message (e.g. Searle) accounts, I am pretty well committed in advance to seeing my own view as Content Sufficient. The challenge for me then will be to provide at least as satisfying an explanation of the small truth about paraphrase as Davidson does, while preserving the straightforward truth evaluability of metaphorical utterance.

1.5. Transparency

The third of my three truths begins with an observation that, when taken in the sense intended, ought to be beyond dispute, but it can at first seem controversial. The observation is this: in the typical case, when we hear a metaphor utterance framed in familiar words, we have an unmediated sense of having understood it—a sense not unlike that when we hear an utterance with familiar words used in a straightforwardly literal way. I shall call this *transparency*, though, for reasons that I will discuss shortly, I cannot say that I am perfectly happy with the term. In any case, however labelled, what can make this claim seem controversial is its use of the word ‘understanding’, so I shall begin by explaining just how little weight I put on it. (Though, as will become clear, this weight is not negligible.)

When Romeo says that Juliet is the sun we are no more brought up short by this than if he had said, for example, ‘I love Juliet’, or ‘Juliet is standing on the balcony’. As is well known, speakers of a language simply do hear its sentences as meaningful, and this is not optional—it is not something that can be turned off. This mandatoriness is rather like that of object recognition: when you switch on a light in a room, you cannot control whether what you see is a table in the room—you just do. It is this immediate ‘getting’, whether of a sentence in a familiar language or a scene, that I call transparency, and my claim is that it is just as true of metaphors as it is of those utterances we regard as unproblematically literal.

Breaking with my self-denying ordinance in respect of examples, it might help make the point about transparency if I ask you to cast your eyes over a few more complex examples:

- (2) Out of the crooked timber from which men are made, nothing straight can ever be built.
- (3) In cities you build a language of circumspection and tact, a thousand little intimations, the nuance that has the shimmer of rubbed bronze.
- (4) Swerving at the last moment to avoid innocent bystanders, his argument came to halt.
- (5) One can see in her prose the rough timbers that a more careful builder would have covered over.
- (6) He has the personality of a traffic cone.²¹

Each of these is certainly more complicated than the simple ‘A is B’ form that typify examples used in many discussions of metaphor, and I hereby promise that, in Chapter 4, I will show how the account offered in Chapter 3 handles these syntactically complex forms. But in the present context the issue isn't complexity but rather a sort of simplicity. My guess is that none of the above is familiar enough (if familiar at all) to have lost any freshness, yet, even though wrenched from their natural contexts, I do not think these sentences give us any more trouble in comprehension than literal sentences of similar syntactic complexity. Of course, I could have given

²¹ Sentence (2) is of course from Kant (1784/1912: 23); (3) is from Don DeLillo's *Underworld* (1999: 446); (6) is a typical journalistic remark. The others were fabricated for the occasion.

examples of metaphors so commonly used as to make certain of transparency, but it would be easy to imagine someone objecting that it was lack of freshness rather than transparency which I had thereby shown.

In saying that we understand Romeo's remark about Juliet and the sun, or that we can take the metaphors in (2) to (6) in our stride, I am not claiming that we understand them immediately in a sense appropriate to, say, literary study or psychological and philosophical investigation. Certainly, if pressed, most of us could say something by way of elucidating them, but that is not what is at issue here. Even if we have nothing enlightening to say, it is still true that we hear these remarks as wholly meaningful sentences in English. Contrast this with sentences in unfamiliar languages, or with sentences containing unfamiliar words and constructions. For example, if you heard:(7) Ich habe ein Fernsehapparat, and spoke no German; or:(8) Henry was scaphocephalic, or, to use the old standby:(9) The horse raced past the barn fell, I would expect you to be brought up short—to find these sentences anything but transparent. Unfamiliar with German, the first would be merely a kind of noise, the second, since it is far from common knowledge that 'scaphocephalic' means 'having a head in the shape of a boat', would certainly be puzzling, and the third example is a particularly good case of a garden-path sentence; no matter how many times we hear it, we are led up that path and away from comprehension. Focus on (8). Our understanding, when we do achieve it, is *mediated* by our having to do some kind of semantic investigation: either we look up the word (not in an abridged dictionary), or work it out from any Greek we remember. My claim is that, with metaphor, it is simply not like this. No one has to do this sort of semantic research on the words in Romeo's utterance, or in the others, and, for all the difficulty we have in saying what metaphors mean, we do understand them, albeit in a suitably lightweight sense, on first hearing.²²

Note that I am not claiming that when we hear Romeo's remark we understand it in such a way as to make philosophical worries about metaphor otiose. Far from it. At this point, I have said nothing about what it is that we understand immediately—what

²² Of course there are metaphors which do startle us and whose uptake is less immediate. But there are literal sentences like this too. A more realistic picture than I have given would describe a spectrum of cases from the more to the less transparent. My point then would be that metaphors would be no more likely to be lower down this scale than literal utterances. This is especially obvious if you focus on everyday metaphors and not simply on those in 'high-octane' poetic contexts. One of the reasons that I find the Romeo example apt is that, though it does occur in a work of fiction, it is typical of a whole range of metaphors anyone might use in descriptions of people and places. (Note: by 'everyday metaphors' I don't mean so-called 'dead' or 'frozen' metaphors like 'the mouth of a river'. In Ch. 4, I shall have a lot to say about these sorts of case.)

it is that we find transparently meaningful—insisting only that our understanding of metaphor matches in transparency our understanding of undeniably literal utterances. As you will see, there are candidates for what is so understood that might explain this without themselves being very revealing about metaphor. But examining the consequences of transparency is work that lies ahead, and any substantial constraint on accounts of metaphor will be based on that work and not merely on the fact that metaphorical utterances share transparency with those that are not metaphorical.

Despite this disclaimer, and despite my putting so little weight on the notion, I do think that ‘understanding’ is the right word to use in describing what happens. I want to rule out as simply false the idea that what we experience on hearing Romeo's remark is merely a sense of familiarity with word-forms and syntax. After all, we recognize these forms in the sentence, ‘Henry is scaphocephalic’, but this sentence is nonetheless, for most hearers, opaque. Thus, it is not enough to say that, on hearing Romeo's remark, what we ‘get’ is that he has used what are presumed to be English words in an easily parsed construction. Instead, though in some not too heavy a sense, we *understand* him. In what sense? This is the question that needs to be explored.

Just before doing so, let me interpolate a note about what lies behind my choice of the label ‘transparency’. In explaining what I take transparency to be, I have at times described our grasp of Romeo's remark as immediate. However, while I once thought that ‘immediacy’ could serve as a label for the third truth, I have come to see this, and even the term ‘mandatory’, as misleading, even though they capture something missed by ‘transparency’. The observation that I am calling on—my third truth about metaphor—is not intended to usurp psycholinguistic research into metaphor comprehension. But talk of immediacy and mandatoriness can suggest just that. The one makes it sound as if speed of uptake is the issue, and the other is redolent of debates in psychology about modularity. The thought these terms suggest is that, in claiming transparency as a fact, I am anticipating—from the notorious armchair—results of psycholinguistic experiments; experiments which aim to investigate empirical questions such as the time-course or degree of modularity of our comprehension of metaphors. There is such research, and aside from controversies about its methodology, it is far from conclusive. Some experiments suggest that we comprehend metaphors just as quickly as we do literal sentences, and some that metaphor comprehension is slower. Still, whatever we come to think about the time-course of comprehension, my claim should not be seen in this context. That is why I have settled on the more neutral ‘transparency’. Though one might think that transparency would be supported only by a result showing that the number of milliseconds needed to comprehend a metaphor is not significantly different from that needed for a literal utterance, this is simply not the case. It is wholly compatible with the transparency point that metaphor comprehension actually lags behind the comprehension of literal utterances. Transparency is a claim about *how* metaphors strike us, not about *how fast* they do so. Neither ‘Juliet is the sun’ nor ‘Juliet is on the balcony’ bring us up short in the way we are brought up

by 'Henry is scaphocephalic' or, if we speak no German, by 'Ich habe ein Fernsehapparat'. And this remains true, even if it turns out that the first of the sentences about Juliet requires, on average, some number of milliseconds more of comprehension time than the second.

With these clarifications, we can now turn to further business of this section: the consequences of accepting, as we should, the truth of transparency. I have been insisting that we are not brought up short by Romeo's remark because there is a sense in which, when we hear it, we understand it—we take it in just as we do literal utterances. I have also insisted that, though there may be issues about the depth of this understanding, it cannot be as shallow as the mere recognition of grammatical form. Against this background, what transparency requires is that any acceptable account of metaphor must put in place something in the realm of meaning or meaningfulness that is made available on hearing metaphorical utterances using familiar words. In short, any acceptable account must have the resources to explain transparency. This might not sound like much to ask, but it turns out to be a surprisingly strong constraint.

An obvious first move, relevant to various accounts, would be that when we hear Romeo's remark, we understand it as asserting of Juliet that she is literally the sun. Many regard this as at least one meaning of Romeo's utterance that is available to us on first hearing, and thus might account for its transparency.²³ Moreover, writers as otherwise diverse as Black, Davidson, and Searle could all sign up to this; even if they do not agree about the correct account of metaphor, they could all appeal to literal meaning to explain transparency.

In spite of these advantages, however, this is not, and cannot be, the whole of what we have immediate access to when we hear (R). To see this, imagine a speaker who unproblematically understands (R) as a literal assertion, but who is, as I shall say, 'metaphor-blind'. Such a speaker would have no trouble in drawing consequences from Romeo's remark—he would insist that Romeo had said that Juliet was a long way away, appears in the eastern sky in the morning, etc.—but would find utterly baffling the natural remarks *we* might make in commenting on (R). Here I have in mind comments such as that Romeo couldn't live without Juliet, that she gives his life point, that she helps him to understand the world. These comments are not to be understood as somehow giving the meaning of Romeo's remark. Nothing I have said commits me to anything as strong as this, either in respect of what is needed for transparency or, come to that, what is ultimately needed for an account of metaphor. The point to hold onto is that these kinds of elucidation come naturally and intelligibly to those not metaphor-blind, to those for whom (R) is transparent *as a metaphor*. When these 'normal' hearers encounter (R), and then hear these sorts of elucidation, they do not find them puzzling: something in their initial encounter with the sentence must have prepared

²³ I doubt that metaphors *in general* have a literal meaning that happens to be false. Instead, I side with those who think this an artefact of looking at subject-predicate metaphors, and that more realistic examples would show metaphors taken literally to be largely nonsense.

them. However, if what we ‘get’ transparently when we hear (R) is simply its literal reading, then it would leave inexplicable our lack of puzzlement when elucidations are advanced.

Here there would seem to be an obvious rejoinder. Perhaps what we get is not literal meaning on its own, but that *plus* something. We hear the literal meaning and, if we are not metaphor-blind, we also recognize that we are not meant to stop there. Spelling this out, we can say that those who are metaphor-sighted take in Romeo's utterance literally, but also take it as including some further encouragement that could be put as: ‘Think of this utterance as metaphorical’. It seems to me that this story is on the right track, but it doesn't yet go far enough.²⁴ Our preparedness to find comments about Romeo's remark intelligible cannot be explained by our simply noting its metaphoricality; the comments are specific to the actual words used in the utterance, they would not have been intelligible if Romeo had said that Juliet was the moon. Any reasonable story about what the metaphor-sighted auditor grasps must appeal to something more focused than the mere identification of an utterance as metaphorical.

Given this, the next move is pretty well forced. Recognizing that our initial grasp of a metaphor is not exhausted either by its literal meaning, or by this together with a general licence to treat the utterance metaphorically, why not make the licence more specific. Instead of: ‘Think of the utterance as metaphorical’, make it: ‘Think of the metaphorical rendition of the particular words used in the utterance.’

This last suggestion seems at first an improvement. Transparency now is explained by our being struck by a literal meaning of certain words at the same time as we are struck by a need to go beyond these very words. But for all that it seems an improvement, it is still too general. The instruction ‘Think of the metaphorical rendition of these particular words’ is not yet pointed enough to explain why our initial grasp of a metaphorical utterance suffices to make intelligible the highly specific elucidations and commentaries that we might come across. Thus, someone could grasp the literal meaning and also the instruction, but, until the instruction is actually carried out—something that would destroy the sense of transparency we have—the connections between what is grasped transparently and elucidations of it would be unexplained.

I can imagine someone struggling a bit with these last points, so it might be a good idea to have in front of us a clear example of a proposal that would in fact do the trick; a proposal that honours transparency while at the same time allowing us to find intelligible the connection between what is initially grasped and the elucidations that typically follow. We can then think of this example as a benchmark: any acceptable account of metaphor would have to do at least as well in its handling of transparency.

²⁴ It might also be wondered whether it goes too far, since it makes the identification of an utterance as metaphorical crucial to our grasping it. Perhaps the truth is that we grasp an utterance which happens to be metaphorical, but manage this without being possessed of an ability to identify it as metaphorical. In any case, since I don't believe this story, I don't want to expend too much effort defending it.

The proposal is based heavily on Davidson's discussion of metaphor, though there are reasons for thinking that he would not accept it. I shall call it the 'Image Account'. According to this account, we who are metaphor-sighted take Romeo's remark, not as asserting some literal falsehood, but as putting an image (not necessarily visual) in front of us. That is, we do not take Romeo to be expressing the thought that Juliet is the sun; in fact, we do not take him to be expressing any thought. Rather, we take him to be in effect producing a not quite sentential utterance of the following sort: (R') Juliet as the sun. Since what I am suggesting here is only a sketch for the purposes of seeing how transparency can be accommodated, and as it is a sketch of a view that I do not think is ultimately correct, I don't want to get too involved in details. Clearly, there are issues about how, if at all, we could extend the Image Account to other forms of metaphorical utterance, and also issues about what triggers the 'reading' of (R) as (R'). But, leaving these issues on one side, what I do want to emphasize here is how little would be necessary for someone to take an utterance of (R) in the way displayed in (R'). With one small difference, the words of both are the same, so there is no reason to worry about the transparency of (R) carrying over to (R'). Indeed, we might even imagine that there is some mechanism which accesses structures of the form (R') whenever those of form (R) are heard, and that this mechanism operates across the board, and not just when metaphors are at issue. I do realize that this speculation violates my promise not to trespass on psycholinguistic research, but it doesn't trespass very far. The idea would be that even when we hear a literal sentence such as: Henry is English, there is a parallel activation of: Henry as English. Thus, not only do we get, *transparently*, the thought expressed (the true-false claim that Henry is English), we are also at the same time made aware of an image—though not obviously a visual one—of Henry as English. In the case where a sentence is intended only literally, our awareness of this second 'reading' is at most muted, and it therefore doesn't play a central or direct role in further conversational exchanges. But the story is different with metaphor: on the Image Account it is the second reading that dominates.

If the Image Account were in fact to work in the way described, then we would have no trouble understanding both the transparency of metaphors and the fact that certain sorts of elucidation are unpuzzling. First consider the issue of transparency. What the Image Account offers us is a simple scheme showing how a metaphor can be wholly present in familiar words of a sentence without our having to make a detour via literal meaning, even if this is available. This neatly explains why metaphors are as transparent to us as literal sentences. We hear the words of the

metaphor and are aware both that they seem to express a thought and at the same time present an image or picture.²⁵ In so far as we do recognize the utterance as metaphorical—and a story about how we do this is not something that need be an intrinsic part of the Image Account—we will suppress the literal reading in favour of the image. But all of this happens downwind of our being confronted with the sentence, and so does not interfere with the sentence's transparency.

Consider next the fact that a certain sort of elucidation strikes us as unpuzzling. This is just what one would expect given that we can be simultaneously aware *both* of a possible thought expressed *and* of an image. Being told in connection with (R), for example, that Juliet is essential to Romeo's life, those who are metaphor-sighted will not be brought up short. Though the remark is scarcely intelligible as a comment on the literal thought that Juliet is the sun, it is perfectly natural to understand it as a comment on the image of Juliet as the sun. Indeed, it might well be that a conversational context in which such comments are made helps to suppress the literal reading, and thereby should be counted as among the ways in which we identify an utterance as metaphorical. Of course, if this speculation were true, it would mean that metaphor identification is something that can take place *after* we have understood (in my transparent sense) what are in fact metaphors.

This last point is by no means something that I am insisting on, but it is intriguing, hinting as it does of a connection between the Image Account and a certain well-studied feature of online utterance interpretation. When one hears the sentence: John found a bug in the room, it seems, perhaps surprisingly, that we access both of the meanings of 'bug', and therefore take in our stride either of the following continuations: He went in search of the insect spray. He hadn't realized that he was the subject of surveillance.²⁶ What I am suggesting is that something like this same model might well apply to metaphorical utterances. We hear (R) and immediately (transparently) access both the thought-expressing version of (R), and the image-presenting version (R'). Because both of these are available to us from the beginning, we are ready to understand continuations that show (R) to be metaphorical; indeed we would be less

²⁵ This leaves out Davidson's suggestion that we (i) recover the literal meaning of the metaphor, (ii) notice that it is absurdly false (e.g. in the case of Romeo and Juliet), and then (iii) entertain the picture of Juliet as the sun, thereby finding something in Romeo's speech act which—though not meaning or content—makes it intelligible. The first two stages of this are, in effect, Davidson's speculation on the mechanism by which we identify an utterance as metaphorical. The Image Account takes into consideration only the third step and builds on it a story about the content of a metaphor that Davidson would probably not accept.

²⁶ It would take too long to explain the careful experiments that have established this. Suffice it to say that experimenters have looked closely at the priming effects of each of the meanings of 'bug' and have found results suggesting that, on hearing the word in the first sentence, subjects are primed to recognize more quickly words appropriate to both meanings. I have hinted at this by saying that we take both of the continuation sentences in our stride, but it is in reality a little more complicated than this.

surprised by an elucidation relevant to the metaphor version than by the continuation: 'Light from her takes more than eight minutes to reach us.'

Leaving aside as mere speculation the mechanism by which the Image Account accommodates the elucidation point, the main thing to keep hold of is that it can accommodate it, and at the same time provide us with a way of thinking of transparency that is simple and plausible. To repeat, on the Image Account we are at least given something that is available to hearers when they encounter a metaphorical utterance—something that is a quite specific and familiar content with respect to which transparency is unsurprising.

That said, I think we must ultimately reject the Image Account, and this because it fails in just the way that Davidson's does: it is simply unable to accommodate truth. On the Image Account, 'Juliet as the sun', is not even a sentence, and hence is certainly not suitable for expressing any kind of truth. But the Image Account certainly shows the way in respect of transparency; it gives us a model to which we should aspire. Can any of the other accounts of metaphor considered earlier match this? That is, do they put in place something that we can plausibly be said to understand when we first encounter a metaphor, thereby explaining transparency? And is what is put in place the right kind of thing to account for the lack of puzzlement that we feel when later on certain elucidations are offered? Perhaps surprisingly, I think the answer is 'no'. Black's account will be the first exhibit, and its handling of transparency will be considered in some detail, as some of the points raised will figure later on. But the criticism actually applies to a whole range of Content Sufficient accounts.

Black's view, at least in the early version, is that the literal meaning of the words in a metaphorical utterance is only a part of the story.²⁷ When we hear Romeo's utterance—and somehow identify it as metaphorical—this starts us off on a search for a second, metaphorical meaning of the predicate 'is the sun'. The full story stresses the importance of the non-metaphorical subject in guiding this search, as well as of course the 'associated commonplaces' that go with the predicate. But the full story is not important here, and the following summary will serve: one begins with words whose meaning is familiar and literal, and, having identified the utterance as metaphorical, one elaborates and extends these meanings in reaching for the metaphorical content of the utterance.

Many questions have been asked about the transition from literal to metaphorical meaning in this sort of account, and disputes over the answers have generated a large literature. But my question is different: I want to know what it is, on Black's account, that we—the metaphor-sighted—understand when we first encounter a typical metaphorical utterance. It cannot be the metaphorical meaning itself, since the latter is the *result*, not the starting point, of a process of reflection. It cannot be simply the literal meaning of the words because, as has been noted, it is perfectly possible for someone to grasp this without having a clue about metaphor. Finally, it

²⁷ If it is felt that I am being unfair to Black and that his considered view is closer to an Alternative Message account like Searle's, this won't matter for present purposes. For Alternative Message accounts fair no better in respect of transparency.

cannot be the literal meaning, together with some general recognition that we have a metaphor on our hands, because this leaves mysterious our finding quite specific elucidations of the metaphor unpuzzling.

As noted above, my worry about Black's account must be distinguished from the more usual ones, but it is importantly related to them. I am not saying that it is mysterious how we ever get from literal to metaphorical meaning, or that we have no model of how the two kinds of meaning fit together. It is true enough that standard models of ambiguity and polysemy do not fit metaphorical meaning very well, and that there is no good explanation in the offing of how it is, as one writer puts it, that the literal is still active in the metaphor.²⁸ But these worries can be deflected by just insisting that metaphorical meaning is at the end of the day special; speakers who are sensitive to metaphor make the transition to metaphorical meaning—a transition different from ordinary ambiguity or polysemy—and do so while keeping a grip on the literal starting point. But my worry cannot be deflected in this way. For what is missing in Black's account is a story about the content of our understanding *before* we make the transition. Without such a story, we must consider the typical speaker's encounter with Romeo's claim as pretty much the same as her encounter with 'Henry is scaphocephalic'. To be sure, the words in Romeo's utterance are all familiar, but, as I have taken some trouble to argue, the familiarity of the word-forms and their literal meanings is deceptive. On their own, they do not suffice as an account of what we grasp on hearing Romeo, and their ultimate transmutation into metaphorical meanings comes too late.

This last point invites an obvious rejoinder. I took some pains to insist that I was not anticipating psycholinguistic studies about the time-course of metaphor comprehension. Given this, a supporter of Black might insist that it is possible, as a matter of empirical fact, that the sub-personal processes of utterance interpretation that go to work when we hear a sentence are really rather fast. To be sure, metaphor interpretation is bound to take longer on average than that required for literal utterances; for the literal version of the utterance must be processed to some extent in order to deliver the metaphorical version. But this doesn't matter if both happen fast enough to guarantee that from the hearer's point of view, metaphors are no less transparent than literal sentences.

While initially plausible, this rejoinder doesn't work. For, as noted, transparency is not about how fast we understand, but is rather about what we understand. On the rejoinder just considered, we are asked to accept three things: first, that as a result of sub-personal processes of interpretation, hearers experience literal utterances as having *transparently* some kind of content; second, that as a result of similar sub-personal processes, hearers also experience metaphorical utterances as also having content; and third, that while the sub-personal processes in respect of metaphor take longer, this time-lag is well below the threshold of our awareness, so transparency of content is typical both of literal and metaphorical utterances.

²⁸ Moran 1997: 254. The full citation, which I fully endorse, is: 'the special dependence of the metaphorical on the literal ... makes the literal meaning of a word ... still "active" in the comprehension of its metaphorical use. We are still in need of an account of this activity.'

The first of these contentions is certainly right, and, as I have already said, I have no problem with the time-differential point in the third, though such evidence as there is suggests it is false. However, the second contention is problematic, not because it is psycholinguistically doubtful—something that cannot be judged philosophically—but because it does not cohere with the other claims. What we are asked to accept is that, on hearing a metaphorical utterance, we process the words in the utterance for their literal meaning, and then go on to re-process them so as to arrive at a metaphorical version. All of this is held to happen sub-personally and quickly enough to give us the sense of the metaphorical version as transparent, as something of which we are unmediatedly aware. But by the first contention, the sub-personal processes of interpretation, when applied to an utterance that has a literal reading, in fact do produce such a reading and, as it were, hand it on to awareness. That is, they yield a reading of which we are unmediatedly or transparently aware. Yet the second contention seems to require that, when faced with the words in a metaphor—words that the account requires to have a literal reading—we are not aware of this literal reading. How does this happen? Do we sub-personally suppress the literal reading, waiting instead for the metaphorical version to emerge into the daylight of transparency? Suggesting something like this is not merely desperate, it is hopeless. It is desperate because it would require a degree of ‘look ahead’ not usually thought appropriate to sub-personal processing. But it is hopeless because the reading that is somehow ignored or suppressed is in fact a *reading*, and as such it is something at the personal level. So, it is not merely that the sub-personal activity must look ahead, it must somehow be able to consult an output at the personal level, and, depending on what is found there, take action. This makes a nonsense of the whole story about the sub-personal which was intended to accommodate transparency within a Black-type Content Sufficient account.

The argument just presented depends heavily on the distinction between the literal and the metaphorical. That some such distinction is commonly made and is intelligible should be beyond doubt, but it is far from obvious whether we can spell it out in a theoretically interesting way. Yet, as I presented it, my argument depends not only there being such a distinction, but that it feature in the Black-style story about utterance processing. However, if that story could be revised in such a way as to dispense with the literal/metaphorical distinction, at least for the purposes of utterance comprehension, this might well allow Black's account to escape my strictures. And there is such a revision.

One begins by thinking of predicate expressions as having associated with them features or properties, and that, encountering some such predicate expression in a sentential context, we have access to these features, they somehow become activated. Thus, in no particular order and with no claim to exhaustiveness, features of ‘is the sun’ might be: is visible in the sky, is a long way away, is at the centre of the solar system,

undergoes nuclear fusion,gives us daylight,is necessary to life,warms us,keeps the Earth in its orbit,is responsible for metering time.Now, while we access these and other features on hearing 'is the sun', in certain specific contexts, we give priority to one or another subset of them. Thus, on hearing the sentence:The cause of recent power failures is the sun,we think of that subset of features which are most appropriate to a nuclear-powered star in our vicinity. We are thus unsurprised when the narrative goes on to speak of sudden bursts of charged particles that can emanate from the sun, and have been known to interfere with power transmission. However, if the context is that set by Romeo's utterance, then a different subset of these features—those having nothing to do with nuclear fusion or distance from the Earth—are given heightened importance, and we should therefore be unsurprised when we encounter continuations which do not depend on purely astronomical features of the sun.

More could be (and will be) said about this proposal, given here only in outline, but the idea should be clear enough. And it should also be clear how this proposal might be thought to get around my previous worry about coherence. For we no longer have to assume that the sub-personal systems responsible for processing utterances have to look ahead to, and perhaps suppress, a literal reading in order to come up with a metaphorical version. Indeed, on this proposal, one doesn't even have to insist that there is a theoretically interesting literal/metaphorical distinction. All that has to happen is that sub-personal mechanisms make an appropriate selection of, or give emphasis to, some subset of features associated with the predicate expression.²⁹ Thus, if it is as a cause of electrical power cuts, features including solar nuclear fusion are stressed. But if it is as a predicate of Juliet that we come across the expression, it is that feature set including warmth, light, and life-giving that is to the fore.

Once some subset of features is singled out sub-personally, our understanding of the whole utterance can come to awareness in a way which strikes us as transparent. Somewhere down the theoretical line, the utterance might come to be described as 'literal' or 'metaphorical', depending now simply on which set of features happens to be associated with the relevant predicate expression, but this distinction doesn't play a role in the processes that give rise to the initial understanding.

²⁹ I keep talking about the 'sub-personal', not because I am trying to anticipate the work of psycholinguistics, or because I endorse such talk, but merely to have a name for whatever might take place before we have that familiar feeling of understanding an utterance that I call transparency.

This proposal is certainly an improvement on the previous one. However, it still doesn't give us what is required for transparency. To see why, focus on the subsets of features which are differentially stressed in comprehension. A first thought here might be that, in speaking of a sub-personal system as being able to make an 'appropriate selection' of one or another subset of associated features, we are crediting such a system with powers it simply couldn't have. This objection falls roughly into the same category as my earlier objection to the less articulated account. There it was argued that a sub-personal system couldn't *coherently* be said to make us aware transparently of the literal sense of an utterance, and at the same time be able to pass over such a literal version, without our being made aware of it, when the relevant utterance happened to be metaphorical. However, the imagined objection to the new proposal is not a claim about coherence, but about plausibility; it is not that we are requiring one and the same sub-personal process both to do and not to do something, it is merely that we find it difficult to imagine how a low-level process could manage something as difficult as making an informed selection from amongst various sets of features. So this first objection is not as conclusive as the earlier one: difficulty is simply not the same thing as incoherence.

It is familiarly unreasonable, in psycholinguistic contexts, to argue from what *seems* difficult for a low-level process to what *is* difficult. For example, it is perfectly possible—and would fit in nicely with Black's interactionist view—for access to the subject term in a metaphorical predication to be the trigger for highlighting appropriate subsets of features. And there is nothing about this kind of access, or its triggering effects, which is so difficult as to be beyond a low-level system. One kind of subject ('The cause of power failures') predisposes the activation of one set of features which are associated with the predicate ('is the sun'); the subject term 'Juliet' predisposes the activation of a different set.

Still, though this first objection to the revised Black-style view is wide of the mark, concentrating on the subsets of associated features will reveal a much more devastating worry. The list of features associated with 'is the sun' have this important and unsurprising characteristic: they are all features of an astronomical body, the star around which our planet orbits. It is easy to lose sight of this. Encountering 'is a source of warmth' in the list of associated features, it is all too easy to think of 'warmth' in what is in effect a metaphorical way. But this would be unjustified, or at least not yet justified. The sun's warmth is no less a physical fact about it than its location or its undergoing nuclear fusion; and the same goes for its being a source of light and for its being responsible for life on earth. None of the features in the original list, nor therefore any subsets of such features, apply without equivocation to human beings. And while the list could be extended, this wouldn't change the point at issue.

Taking this on board has an immediate consequence: we cannot generate from the list of such features a rendition of 'Juliet is the sun' which would explain, what seems obviously true, that we take in this sentence transparently, and do so in a manner fitted to its undoubted metaphorical nature. No doubt some of the features in the list will more easily lend themselves to transference from the stellar to the human context than others. This is what lies behind the suggestion that focusing on

certain subsets of features might help with transparency. Yet the plain fact is that even the most transferable features are just not true of the metaphor subject.

Whether or not we call it a literal/metaphorical distinction, there is no getting away from the fact that expressions in metaphorical contexts have some kind of primary (before transfer) and secondary (after transfer) meanings or, perhaps more accurately, associated features.³⁰ While the idea of such movement from literal to metaphorical, or primary to secondary, might be perfectly reasonable in some larger picture, what this *movement* suggests is that metaphorical utterances—as metaphorical—should not be transparent to us. But, when due account is taken of what transparency requires, they most certainly are transparent. What Content Sufficient accounts like Black's lack—what any account which posits literal meanings or primary features lacks—is something appropriate to put into play when we initially encounter metaphorical utterances. The Image Account shows the way: one hears a metaphorical utterance and comes away with an image rather than with any thought expressed by the literal meaning, or by some subset of associated features. This image is of course painted by the literal or primary meanings of the words in the utterance, but the image itself is distinct from any literal or primary *judgement* that the utterance might be taken to express. It may well be that an initial thought of Juliet sharing the feature of warmth-giving with the sun sets us on the path of understanding. But since Juliet no more shares this feature with the sun than she shares the feature of undergoing nuclear fusion, it is wholly mysterious why metaphorical utterances strike us transparently.

The basic problem that Black-style Content Sufficient accounts have with transparency is even more starkly problematic for Alternative Message accounts. This is because the most straightforward of these accounts credits hearers with some initial interpretation of a metaphorical utterance that is most certainly not itself metaphorical. For example, Searle would insist that we understand Romeo's words as saying that Juliet is in fact the sun, and then we move on from this to what is in the end a metaphorical reading. As he says, metaphorical utterances literally say that S is P (though in most cases S is certainly *not* P) and we come to understand the speaker only when we appreciate that he is in fact saying that S is R, where the transition from one assertion to the other is governed by a fairly loose network of principles of interpretation. Whatever else can be said for or against this, it makes transparency wholly mysterious. Worse than having no candidate rendition to serve the needs of transparency, Searle's account has one that is precisely wrong; saying that what we grasp is literal meaning, or any of its variations, makes it not merely difficult but impossible to accommodate transparency.³¹

³⁰ I do accept that the two pairs of distinction are different. However, when it comes to transparency they cause the same kind of trouble. See Kittay 1987 for a discussion of the two distinctions.

³¹ There are more subtle Alternative Message accounts than Searle's, and the issue of transparency is less clearly troublesome for them. In particular, the account offered by Fogelin (1988) does suggest a way, somewhat along the lines of the Image Account, in which transparency might be handled. However, I reserve a more detailed discussion of Fogelin for Ch. 5 because his view raises issues that are best discussed along with certain other rivals to my account rather than on its own.

Let me conclude with one final point about transparency which might help to tie up a few loose ends. Many of the examples of metaphorical utterances that philosophers use tend to be on the 'low-octane' side. Thus, Richard is a lion, Jeremy is a rock, Mary is a bulldozer, are typical. However, the point I want to make here about these examples concerns transparency, not the issue of their metaphoricality.³² For whatever else you can say about the above, they are certainly transparent: speakers and hearers of English would be expected to comprehend them with no more difficulty than any straightforwardly literal sentence of the same complexity. Nor is this surprising, given the familiarity of the non-literal versions of their predicate expressions, a familiarity supported by many dictionaries.³³ If, as seems the case, the philosophical diet of metaphorical examples is roughly of the kind shown above, then one can understand why transparency has never seemed much of a problem. Even though many philosophical writers are sensitive to the fact that their examples tend to be low-octane, it is all too easy for them to imagine that whatever is going on in these cases can be extended in some way when the octane level is higher. This doesn't mean that they regard this extension as easy, but once having begun with these sorts of examples, it is difficult to appreciate the issues that transparency raises. For example, if you take: Mary is a bulldozer, to be your initial example of a metaphor, then you wouldn't notice that there is any special problem about transparency; the above is as immediately comprehensible as any literal sentence. However, when, or if, you come to consider richer cases, even only mildly richer ones like Romeo's, you will assume that, to the extent it is transparent, this needs no special further explanation. Perhaps, 'bulldozer' will have a metaphorical entry in the dictionary, and 'sun' will not; perhaps the one is conventional and therefore familiar, the other less so. But this can seem a trivial matter, since you will imagine that hearers can, as it were, make up a dictionary entry for themselves on the fly. The trouble is, though, that by the time you have reached this point, you will not have noticed that our initial grasp of even quite subtle metaphors needs an explanation that cannot be provided by any such process of extension from low-octane examples. The latter are, for the purpose at hand, just like literal

³² It might be argued that these are simply too trite even to be counted as metaphorical; the thought might be that they are simply dead metaphors, and, as some think, to be a dead metaphor is to be a metaphor no longer. I don't think that this is the right view to take about dead metaphors, but discussion of this won't come until Ch. 4.

³³ Interestingly, 'le bulldozer' was introduced into French in 1948 by ministerial decree 'for the enrichment of the French language' but, as far as I can tell is used in contemporary French exclusively in a non-literal way. There may be lessons in this for our idea of dead metaphor, but, as noted, discussion of this comes later on.

sentences; the subtle cases are not: there is nothing available to explain why we find even high-octane metaphors transparent.

1.6. The Shape of Things to Come

Much of the work of this chapter has been negative, or at least appears to be negative. I began by describing what I take to be three important truths about metaphor:

- (i) Metaphors *can* be fully assertoric, and have truth-evaluable content for which, like assertions generally, the speaker is responsible. (Even this softened claim is of course controversial and I didn't argue for it in any detail, but I am not alone in thinking it non-negotiable.)
- (ii) It is inappropriate, and not merely difficult, to try to paraphrase metaphors.
- (iii) Metaphors are as transparent as literal utterances (of similar grammatical complexity).

I then argued that certain accounts fail to accommodate one or more of them. Straightforward Content Sufficient accounts which seek some special or additional metaphorical meaning attached to expressions in metaphorical utterances have an uphill task accounting for the inappropriateness of paraphrase and fact of transparency. Alternative Message accounts certainly fall foul of the strictures on paraphrase, and, in the version Searle champions, cannot account for transparency. A Davidsonian No Message account can be constructed which offers us a clear vindication of the paraphrase and transparency points, though it fails pretty spectacularly to deal with the possibility of truth in metaphor.

Is this pattern of success and failure directly mandated by my classification? That is, is it to be expected that any Content Sufficient or Alternative Message account, merely because they count as such, is unable to deal with the paraphrase and transparency points, and that any No Message account will lack the resources to deal with truth? If I am right in claiming my classification as exhaustive, it is unsurprising that the answer I am required to give is 'no'. For if it were 'yes', then there would be no hope for the account that I develop in the next two chapters.

It may be unsurprising that I think this, but it should nonetheless be surprising that all three truths can in fact be accommodated by a single account. This is because the first seems in tension with the second: it is not at all obvious how metaphors could have whatever content is necessary for truth, while at the same time being beyond even the possibility of paraphrase. In any case, I certainly hope you are sceptical about the possibility of joint satisfaction of (i)–(iii) because it should make you that much more interested in seeing whether, and how, the trick can be pulled off. After all, as I noted at the beginning of the chapter, in trying to find fault with other accounts, I was not merely intending to clear a space for mine, but aiming to clear a space of a certain shape. Having to accommodate the three truths about metaphor, especially given the tensions amongst them, certainly fulfils that brief.

My account of what is going on in metaphor will not emerge until Chapter 3, and even then it will need the support of discussions in Chapters 4 and 5. But, even

though I insisted in the Introduction that this gradualist and minimalist approach is necessary, I don't want to be accused of teasing the reader. So, without any pretence of giving a full-blown account, I conclude this chapter with a sketch of what are the main elements of my view. At the very least, this sketch will guide a reading of the chapters which follow.

In order to find Romeo's utterance intelligible, though without negotiating away the first of the truths discussed in this chapter, we are going to have to find in it a content which is rightly thought of both as down to Romeo—as a content he intended to be found there—and as a content that can be evaluated for truth. With this as an opening gambit, the next moves seem forced. Given that, in (R), Romeo uses words in a natural language, it is difficult to avoid thinking that any content we find is produced in some way by these words or by his use of them. There are basically two familiar ways, each well-studied, this might come about. Either the words in Romeo's utterance have, or come to have, what are usually called 'meanings', and these in effect produce the necessary content; or those same words lead us to the relevant content by indirect means, by our coming to see in the act of utterance some kind of indirect speech act. However, as I have argued, both direct and indirect means of production fall foul of the truths about paraphrase and transparency, and what is needed is something, as the saying goes, completely different.

My account accepts that Romeo's words bring the relevant kind of content to our attention, but denies that such content is in fact *produced* by those words, whether by familiar direct or indirect means. When Romeo says that Juliet is the sun, part of what he does is something almost unnoticeably simple: he uses words in an ordinary way to bring the attention of the audience to an object, the sun. My contention will be that, when we think hard enough and carefully, we will come to see that objects can take on functions that are ordinarily thought of as centrally within the remit of words. One such function is referential: we find nothing problematic in thinking of objects as taking on this kind of function, one we also attribute to words of various categories in natural language. But another function, or so I shall argue in Chapter 2, is in effect predicational, and when Romeo makes us attend to the sun, what makes his utterance intelligible is precisely the fact that the sun can count as a sort of predicate of Juliet. No need for us to find some special substitute meaning for 'is the sun', or some substitute speech act for saying literally that Juliet is the sun. We must simply think through the idea that the sun itself comes to have the semantic properties requisite for the task.

That objects can and, more importantly, do function as predicates is not familiar enough to be controversial. So a substantial part of my task will be, first, getting you to appreciate that there is a semantic function, independent of reference and not essentially tied to words, which nonetheless is recognizably predicational; and, second, that objects (as you will see, in a broad sense of the term) do take on this function. But another important piece of business consists in getting you to resist a natural, but damaging, temptation.

I am scarcely the only one to have noticed that Romeo's utterance makes some kind of reference to an object. But reference to objects in this and other metaphors

counts for most writers only as a means to showing how the words that Romeo uses become intelligible, as either a direct or indirect way of saying something about Juliet. Without naming names at this point, the basic idea is that, in those cases where a metaphor makes an object prominent, the function of the *object* is to set us on the right course for understanding the *words* in the original. Juliet has properties, so does the sun. According to a common idea, when we get Romeo's utterance right, it is because we have managed to find those properties of the sun which match those of Juliet (or are intended by Romeo to match Juliet). It is these properties which then constitute the Content of 'is the sun' in the context of (R), or which give us the means to find some Alternative Message that Romeo intends. Treating the relevant matching of properties as tantamount to a judgement of similarity, we can say that we count (R) intelligible in virtue of a similarity we find between Juliet and the sun.

My discussion of similarity is spread throughout the book because appeal to this notion in connection with metaphor is made in so many subtly different ways (and not always explicitly) that any attempt to mount a single definitive counter-argument couldn't have succeeded. However, though my arguments might have to be sensitive to context, their conclusion is basically the same. When taken *in the right way*, it is certainly true, and not trivially so, that, for example, Juliet is similar to the sun. But before we can understand what that 'right way' is, we have to have in place an account of metaphor. Hence, appeal to similarity in all its guises, while not wrong, is of no use in helping us understand the phenomenon of metaphor. My account suggests why judgements of similarity (of the right kind) appropriate to metaphors are true, but that is the only good word I shall have for similarity and its offspring. Moreover, it is not simply in connection with metaphor that similarity lets us down. There will be places in my discussion of predication where I am sure you will be overwhelmingly tempted to think: but his idea is in essence an appeal to a kind of sharing of properties, to similarity. Resist this temptation. What I shall say about predication is intended as radically alternative to accounts in terms of property-sharing and similarity, and even if you don't accept my account, the least you can do is to see just how different it is.

2

Object and Word

2.1. Introduction

It is a fact, but one not generally thought important, that words, while they may stand for or refer to objects, are themselves a kind of object. This is most clear when we think of words as marks on paper, but it is no less true when words manifest themselves in other ways. Sounds and sign-language gestures are less concrete than ink marks, but they would surely count as particulars in any reasonable ontology. Of course, where words are concerned we tend to be less interested in tokens than types: when searching for a word, any one of its instances will do, and only certain specialists—for example, phonologists, graphologists, or forgers—take an interest in words as particular objects. In the normal run of cases, word tokens, like parts on a production line, are objects designed to be interchangeable, as well as easily manufactured. But such effortless interchangeability shouldn't obscure the fact that individual words, in the token sense, are individual objects.

Thinking about the fact that words are objects might lead one to wonder whether objects could be words. Putting it less mysteriously: it might make one wonder whether objects that are not in any ordinary sense words could nonetheless function in something like the way words do. In this chapter, I shall do more than wonder about this. What I want to explore here is what seems to me the evident capacity of objects to take on roles more usually associated with words. However, even before I begin, two cautionary notes.

First, the notion of an object should not be thought of in too narrow a way. Thus, while the word 'object' is bound to call to mind chairs and tables and other concrete particulars, I encourage you to extend it as well to actions, events, states of affairs, circumstances, facts and the like. These things are, after all, *things*, and it is this non-committal use of the word 'things' that you should take as the model for my talk of objects.

Why do I not then use 'thing' in place of 'object'? There are two reasons for this. In many contexts, 'thing' would no less misleadingly suggest middle-sized dry goods than does the word 'object', so it would be no improvement on it. Additionally, in other contexts, 'thing' is simply not substantiating enough: it can play a purely grammatical role something like 'it' in 'It is raining'. Think here of sentences like: 'Things are bad at work'. While I do not want a word that commits us to a substantiality as of dry goods, I do want some substantiality. For want of a better alternative then, 'object', taken in conjunction with this note, will have to do.

Second, though what I have to say in this chapter will be crucial to my account of metaphor, it constitutes only a first step in the enterprise. As described in the Introduction, the phenomenon of metaphor creates problems for our theorizing about language. These problems arise because a kind of use (or meaning) of certain words in natural languages does not fit comfortably into the theoretical accounts of meaning (and use) currently on offer. However, the primary focus of this chapter is not on the meaning or use of words in natural language. Rather, to repeat what I said above, its aim is to explore the ways in which objects (in the broadest sense of this notion) can take on the functions of words. Of course, in pursuing this subject, I will inevitably have to make reference to the kinds of thing that words do. How else can I suggest ways in which objects might do the same? But though metaphor is among the things that words do, I will not claim that there is an immediate connection between what I say about the linguistic functions of objects and the metaphorical function of words. That will come later.

2.2. Objects as Words

Goodman, probably more than anyone else, has explored what we can think of as the semantic functions of objects and, though his work has had its primary impact in aesthetics, it has wider importance. The obvious place to start is with his notion of exemplification, since this notion is at the focus of his efforts to understand how objects can take on semantic functions. Here is an example that he offers in *The Languages of Art*: ‘Consider a tailor's booklet of small swatches of cloth. These function as samples, as symbols exemplifying certain properties. But a swatch does not exemplify all its properties; it is a sample of colour, weave, texture, and pattern, but not of size, shape, or absolute weight or value’ (Goodman 1976: 53). Goodman says that the swatch exemplifies certain *properties*, but he goes on to describe the relationship of exemplification in terms of predicates, that is, words. If you think of the predicates that the tailor could have used to describe the fabric, colour, pattern, etc. of various bolts of cloth, then each swatch is said to exemplify the relevant description when, on the one hand, the predicates are true of the swatch and, on the other hand, the swatch refers to, or symbolizes them.

In what follows, I shall not directly address Goodman's well-known nominalist scruples. For my purposes, it will simply not matter whether we speak of exemplification of properties, or of predicates and labels. My only interest is in the *kind* of linguistic function that exemplification bestows on objects. (That said, nearer the end of this chapter, I will return to issues revolving around the contrast between predicates and properties.)

Goodman regards an exemplifying object as performing what is fundamentally a referential function, but clearly enough there is more to exemplification than straightforward reference. Here is a passage from *Of Mind and Other Matters*, that addresses this issue head-on: Exemplification, far from being a variety of denotation, runs in the opposite direction, not from the label to what the label applies to but from something a label applies to back to the

label ... Exemplification indeed involves denotation, by inversion, yet it cannot be equated with the converse of denotation; for exemplification is selective, obtaining only between the symbol and some but not others of the labels denoting it. (Goodman 1984: 59)

Goodman here insists that exemplification, while it is referential, should be distinguished from what he regards as a central referential relation, namely denotation. In a case of straightforward denotation, we have the use of a label to pick out some object, but in exemplification an object that is in fact labelled has itself the potential *in some circumstances* to denote the original label. The referential credentials of exemplification come from its being a kind of inversion of denotation but, as he says, we cannot define exemplification as simply the converse of denotation.³⁴ This is because not all labels are exemplified by the things they label, so there will be instances where the converse of denotation will not result in anything we could plausibly call exemplification.

Now, given this, one might hope that there would be some principled way to sort cases of labelled objects into those which do, and those which do not, count also as exemplifiers of their labels—cases in which the converse of denotation really does produce exemplification. For perhaps in this way we could say more accurately and precisely which species of reference exemplification falls under. However, this hope is, as far as I know, unfulfilled in any of Goodman's writings, and he may well have thought that there could be no principled way of accomplishing it. Still, for reasons that will become clear, it is worth taking a closer look at some examples, and even perhaps trying to tease out, if not some fully explicit principle of sorting, then some informative features underlying cases of genuine exemplification. Consider first this predicative label: Member of the Birkbeck College Philosophy Department.³⁵ This expression could be said to denote me—it could be seen as a way of labelling me—though I am sure that Goodman would not count this as a case in which the thing labelled (that is, me) is an exemplifier of the label. Though I am indeed a member of the department, there is nothing about me that encourages the thought that I *exemplify* such membership (take my word for it). For a contrasting example, consider this second predicative label: Cloth of beige and brown woollen plaid. Here, not only does this expression point to one of the swatches in the tailor's book, we have it on Goodman's authority that the swatch counts also as exemplifying the predicate.

While it is easy enough to see that there is a difference between these two cases, it is not easy to make this difference explicit, much less to generalize it into a

³⁴ He wasn't always sure of this: see the admission of a mistake in Goodman 1984: 82.

³⁵ I speak here of 'predicative labels' when what is in question is an expression that depends for its referential success on the satisfaction of its predicates. In so doing, I simply pass over the raging arguments about whether such labels are, in the end, genuinely referential. However important they are in our theorizing in logic and philosophy of language, these arguments are simply not important here.

constraint on exemplification. One might be first tempted to think that the difference turns on whether the object counts as a particularly good or paradigm case of the label. I am a member of my department, but it really doesn't make much sense to speak of me as a particularly good, or paradigm, member of the Philosophy Department. However, while this is true of some cases, as a general criterion of exemplification, it won't do. It doesn't even fit the case of the tailor's swatch, which is certainly an example of the relevant cloth, but is not in any sense a particularly good example, or a paradigm case. Surely, any other sample which fits the label would do as well.

Nonetheless, without insisting that it is a fully general account, there is something informative that can be said about the difference between exemplifying and non-exemplifying objects. Moreover, it is something that is implicit in various of Goodman's own examples. For instance, he says: I may answer your question about the colour of my house by showing a sample rather than by uttering a predicate; or I may merely describe the location of the appropriate sample on a colour card you have. In the latter case, the chain of reference runs down from a verbal label to an instance denoted and then up to another label (or feature) exemplified. (Goodman 1984: 62)

In effect, what Goodman suggests here is that an exemplifying object is one which not merely refers us back to its label, it takes on a further *predicational* task. In referring back to its predicate-label, the exemplifying object reveals something which in some way complements, enhances, or even completes its label. The colour sample or the tailor's swatch do not simply count as instances of the predicate; in both cases, they make it possible for someone to appreciate (or appreciate more fully) what it is to be a certain colour or fabric. One might say that they add an epistemic route to the predicate: not relying simply on our mastery of language, we can use our eyes in deciding whether the predicate applies. In contrast, there is no sense in which I, though counting as an instance of the label: Member of the Birkbeck Philosophy Department, would in any way offer someone a further epistemic route to the general application conditions of this description.

The swatch and colour sample bring our capacity to visualize to bear on the relevant predicate-labels, and in this sense, they make a contribution to the application conditions, or, perhaps more accurately, the application of the predicate. Other cases of exemplification can exploit modes other than the visual; one can easily imagine cases in which sounds or things touched enhance the predicates that the sounding or touched objects exemplify. Were my interest solely in Goodman's views, then this idea of predicate-enhancement would certainly seem to be a useful first step; as noted above, it suggests some principled way to pick out genuine cases of exemplification. But I am after ways in which objects might take on linguistic functions that are rather more general and certainly more radical than exemplification.

What Goodman's discussion reveals, albeit indirectly, is that objects can have a special referential role that he calls exemplification, and that at least part of what makes it special is that objects playing that referential role also fulfil a sort of predicative role. In effect, exemplification is a case in which the referential and predicative roles of objects are harnessed together. But I want now to consider the possibility of objects playing these roles separately.

Reference is easy: no one doubts that, in a perfectly straightforward way, objects can be used referentially. Think of the way in which you might use objects on the table at a dinner party to tell a story about a car accident you have had. 'This salt cellar is my car, the pepper grinder is the bus ...' Indeed, it was just this kind of use of objects that so impressed Wittgenstein.³⁶

While we are familiar with the referential use of objects, what is scarcely ever considered, and is in fact even difficult to bring into focus, is the question of whether objects can be used predicationally. Goodman's story about exemplification suggests that they can be used in *aid* of predication. That is why a brief consideration of exemplification served as my introduction. But now we have to see whether they can be used, not merely as 'sidekicks' to linguistic predication, but as predicates in their own right.

I suspect that this will strike many as an odd topic of investigation. However, I believe that the reason for the oddness is directly traceable to a certain carelessness on the part of philosophers in the differential treatments they give to reference and predication. So, before attempting to show ways in which objects on their own can fulfil a predicational role, I will have to begin further back. What I shall say will turn on certain issues in philosophical logic that have a long history, but I shall do my best to touch on them lightly. As I hope you will come to appreciate, the discussion in the sections which follow is not a digression but is in fact necessary. Without it, my view about objects as predicates wouldn't even be so much as visible.

2.3. The Basic Combination

Anything that we can think of as apt for making an assertion is bound to have at least two ingredients: some device for making clear what is being spoken about, and some device for speaking about it. These two ingredients and the functions they call on are evidently distinct, a fact which has been emphasized in different ways by various writers: When verbs are mingled with nouns then the words fit together and the simplest combination of them constitutes language ... When anyone says '[such and such a] man learns' ... he not only names but achieves something by connecting verb (*rhema*) with noun (*onoma*). (From Plato's *Sophist* as cited by Wiggins 1984: 323)

³⁶ Wittgenstein 1914–16/1969: 7, where he mentions the use of objects in a Paris courtroom as a way of representing a motor-car accident (also see editorial note on p. 7). It is clear that this example played a role in Wittgenstein's so-called 'picture' account in the *Tractatus*, and there is a lot more to this account than the use of objects as referential devices. There will be a note about this later on.

Predicates are not names; predicates are the other parties in predication. (Quine 1970: 27–8) Names name, predicates describe, and, having these complementary functions, names and predicates are made for one another. (Wiggins 1984: 323) These claims about what Quine, and then following him, Strawson, called the ‘basic combination’ are perhaps too obvious to need re-stating. But, obviously true though they are, there is a sense in which part of the simple message they convey has not been taken seriously enough.

In the opening chapter of *Subject and Predicate in Logic and Grammar* (‘The Basic Combination’), Strawson describes various formal asymmetries that obtain between subject and predicate expressions in natural language sentences. For example, he notes ‘we can coherently enrich our logic with what we may call negative ... predicate terms, but there is no strictly parallel way in which we can coherently enrich our logic with negative ... subject-terms’ (Strawson 1974: 6). However, not satisfied with leaving things at this merely formal level, he sets about trying to find those ‘fundamental features of thought about the world’ which underlie the subject/predicate distinction. He summarizes the result of this search as follows: We have seen that our thought about the world involves, at a level which, if not the most primitive of all, is yet primitive enough, the duality of spatio-temporal particular and general concept. To bring these two dualities together [viz. subject/predicate and spatio-temporal particular/general concept] may illustrate both. ... Any sort of judgment to the effect that a certain spatio-temporal particular ... exemplifies a certain general concept, or—equivalently—to the effect that a certain general concept has application in the case of a certain spatio-temporal particular ... would be a judgment of this sort. So the hypothesis suggests itself that the basic subject-predicate sentences are sentences apt for the voicing of such judgments. (Strawson 1974: 20) At the level of a relevantly primitive thought about the world, particular and general concept unite in such a way as to allow the judgement that the particular exemplifies the concept. The basic combination is the device at the level of language apt for expressing such a thought about the world, and it manages this by pairing sentence-parts with appropriate elements of our thought about the world. A subject term is assigned the task of picking out some spatio-temporal particular. Its job is thus referential. Predicates, in apparent contrast, are allotted the task of specifying general concepts.³⁷ Why do I say ‘apparent contrast’? Well, implicitly in this passage, and wholly explicitly elsewhere, Strawson makes it clear that the idea of *specification* is less a contrast with, than a type of, reference. I will have more to say about this,

³⁷ Strictly, Strawson speaks, not of ‘predicates’ as having this function, but instead of it falling to ‘concept-words’ to specify concepts. Thus, for Strawson, ‘man’ is a concept-word, and he reserves ‘predicate’ for expressions such as ‘is a man’—expressions which, anticipating what will be discussed shortly, contain something in addition to the concept-specifying word. Others do not make such a sharp terminological distinction: as will be noted Wiggins is happy enough to speak of ‘man’ as a predicate (and many others keep him company here). This terminological issue will not affect my discussion, but it seemed best to note it at the outset. In the paragraphs that follow, I intend ‘predicate’ mostly in its concept-specifying role, i.e. ‘man’ will count as a predicate, though, when confusion might result, I shall sometimes speak of ‘concept-words’.

since it is fundamental to the point I shall come to make, but some mention must be made first of an element that is so far missing from the picture.

As was recognized in different ways (and in slightly different contexts) by Frege, Russell, and Wittgenstein, Strawson notes that *two* linguistic items, each with the job of referring to, or specifying, some feature of our thought about the world, do not *by themselves* constitute a structure apt for expressing thoughts. Given the story described in the passage above, the linguistic structure: (BC-) Subject (e.g. 'Socrates') 1 Predicate (e.g. 'man'), is little more than a list. The subject refers to a spatio-temporal particular and the predicate refers to ('specifies') a general concept. Of course, the latter is certainly not something that Strawson intends us to think of as a particular—he considers general concepts as somehow different from particulars—but it is a *something*. And the mere juxtaposition in (BC-) of a particular and a general concept is thus not enough to count as expressing a thought or judgement. (The minus sign in the label 'BC-' is intended to mark the fact that what we have there falls short in some way of what is necessary to the basic combination.)

Strawson's remedy here is to add a third to the pair of tasks so far used to characterize the basic combination. He writes: 'there is something contained in the sentence as a whole, or something about the mode of combinations of the aforementioned expressions [particular and concept-specifying expressions], which shows that we have a propositional combination of the first pair of expressions' (Strawson 1974: 22). This third task is one which, when it is realized, transforms (BC-) into a genuine basic combination. Strawson is quite clear that what needs to be marked is the *exemplification* of the general concept by the relevant particular (or, equivalently, is the *application* of the general concept to the particular) but he rightly resists expressing such exemplification by some concept-specifying element, since, to do so, would invite an unacceptable regress. Putting a particular-specifying expression ('Socrates') together with a concept-specifying expression ('man') and another concept-specifying expression ('the concept of exemplification') would get us no nearer what is needed in the basic combination; where before we had two things sitting inertly side by side, this new structure would simply give us three such elements remaining no less inert.

Strawson goes on to spell out how this third function can be realized without inviting the regress. In brief, the suggestion is that, in the simplest kind of case, we count the copula (as in 'Socrates is a man') or the finite verb ending (as in 'Socrates swims') as showing propositional combination, that is, as revealing the exemplification that is there in the primitive thought. In so doing, he maintains, we transform what was the inert listing of a particular and a concept into a structure apt for expressing a truth. However, leaving Strawson for the moment, I should like here to interpose a discussion of these matters found in Wiggins 1984. Wiggins's proposal for handling this third element is pretty much the same as Strawson's, but it is offered in the context of a discussion of Frege, and this will prove important.

Unlike Strawson's, Frege's own picture of the basic combination does not include any element treating of the separable relation of exemplification—a relation that, on pain of regress, calls for some kind of special treatment.³⁸ Instead, Frege suggests a way of grounding the basic combination that simply bypasses both the inertness and the regress problems. Famously, if not lucidly, Frege held that concepts were 'unsaturated'. He regarded it as simply true that they are in their nature hungry for, or, more prosaically, have a gap for, particulars, and are just incomplete until the gap is appropriately filled. There being no separate notion of exemplification at the level of our primitive thought about the world—merely particulars and general concepts—all we need to express any such primitive thought is a specification of the particular and the concept. When the most basic subject/predicate (i.e. linguistic) structure is used to express some primitive thought or judgement, even though the subject and the predicate terms do nothing more than pick out, respectively, a particular and a general concept, there is an affinity between the latter two which, as it were, creates a possible fact or state of affairs. Appealing to a chemical metaphor, it is like the bringing together of a radical and its appropriate chemical complement: once this is done, there is no need of some third functional element to transform the combination into a complete atom or molecule.³⁹

This Fregean conception of the basic combination is not without its problems (to put it mildly). On its own, and leaving aside the metaphoricality of the notion of unsaturatedness, it has a certain consistency. But when combined with various other Fregean constraints on concepts, and on the specification of concepts by predicates, a well-documented tension, even incoherence, emerges. Wiggins 1984 offers a fix, intended to prevent any real cracks from opening up in the first place. However, what I shall suggest is that this fix leaves behind an important insight of Frege's. Moreover, though Frege never carried through on it, this insight gestures in the direction of a point about predication that I think is crucial in general, and certainly crucial for the central thesis of this book. First, a brief account of Wiggins's efforts on behalf of Frege.

The fix begins with the following stipulation: 'entities like *man*, *horse* ... such entities—let us call them concepts—are not objects and they are neither saturated nor unsaturated' (Wiggins 1984: 318). Wiggins's aim here is to enlarge the Fregean picture. Whereas Frege distinguished between saturated items (objects) and unsaturated items (concepts), Wiggins adds a third category, namely items such as *man*, *horse*, etc. which are neither saturated nor unsaturated, and he proposes, on Frege's behalf, that this new category is a better claimant to the title 'concept'. Of course, if we accept this alteration to the Fregean picture, we shall no longer be able to say

³⁸ Frege's views are familiar, and I don't aim here for word-by-word exegesis, but the main outline of his account of these matters is in 'Function and Concept' and 'On Concept and Object' in Frege 1952.

³⁹ At one point, Wiggins asks (rhetorically): 'How is it that he who mentions something complete and then something incomplete thereby gets to say something?' (Wiggins 1984: 324). But this is not really fair to Frege. Rather, we should say: he who mentions something complete and then *something which completes it* has certainly brought our attention to something in the world which is a unity. And one can then put in Frege's mouth the counter-rhetorical question: if we bring a completed state of affairs to public notice what more need be done to count as saying something?

that juxtapositions of objects and concepts suffice as the ground of the basic combination. On Wiggins's new conception of a concept, these pairings are simply inert. Given this, we now have a need, which wasn't there in Frege's original account, for something to animate the pairings so that, when reflected in language, we can see them as the bases of the propositional combination. Wiggins addresses this need as follows: 'What the copula does on this alternative view is to *combine* with a concept-word or predicate to produce an unsaturated expression that will in its turn combine in the fashion Frege himself describes with a saturated expression to produce a complete sentence.'⁴⁰ This is a surprising move. Frege's treatment of the difference between concepts and objects was unequivocally ontological: the former are saturated *entities*, the latter are not. In the first passage above, Wiggins's stipulation of a new category of object—the neither saturated, nor unsaturated—fits in well enough with Frege's ontological leanings. But then in the second passage, one which follows on directly from the first, there is a sea change. What are there described as saturated or unsaturated are no longer entities, but rather expressions.

Wiggins certainly seems aware of doubts that might be raised about the transposition from entities to expressions. He writes: Finally, is nothing left at all of Frege's view that the concept is something essentially predicative? One thing remains at least. Our enrichment of the primitive categorial basis to allow expressions of category B [i.e. concept-words such as 'man'] to be coordinate with expressions of the category N [i.e. *subject-terms or names*] permits sentences of both forms 'expression *e* refers to [N]' and 'expression *e* refers to [B]', but it preserves a Fregean asymmetry. At least in the semantics we have given, the copula always combines with the predicate word, not the singular term, to discharge the predicative function in the sentence. (Wiggins 1984: 328) But this response, while it begins with the right question about the move away from ontology, fails to give the right answer. Wiggins's repair manages to give concept-words some ontological work: it allows expressions such as 'man' to refer to entities, in this case the concept *man*. However, when it comes to the Fregean asymmetry between concepts and the referents of subject-terms such as 'Socrates'—the asymmetry between saturated and unsaturated entities—everything happens 'in the semantics'. So, the right answer to the question which opens the above passage should be: nothing at all is left of a concept as some *thing* that is essentially predicative (unsaturated).

I am pressing the importance of Frege's original ontological asymmetry because it is crucial to understanding what I take to be his fundamental insight into predication. Yet, independently of faithfulness to Frege, there is an internal problem with the move from ontology to language that will also prove relevant to that insight, and which I must therefore discuss first. However, since Wiggins says little more relevant

⁴⁰ Wiggins 1984: 318. Note that Wiggins is clear that the copula is not the only thing that does the trick, since of course finite verbs and various other forms can serve as well to express propositional combination. However, so as not to clutter up the text, I will follow his lead in allowing talk of the copula to stand in for other grammatical constructions.

to this issue than is contained in the above passage, I can best show what this problem is by returning to Strawson's account. As will be seen, Strawson saw and at least attempted to address just this problem, and did so within what is pretty much the same framework as that later proposed by Wiggins.

Strawson briefly considers and rejects what he describes as the Fregean idea 'that the function of concept-specification really included what has just been represented as the distinguishable function of indicating propositional combination'. He writes: But we here do not understand the word 'concept', or the function of concept-specification, in any such sense. Rather we understand them in such a way that, e.g., the concept green can equally well be said to be specified in each of the three following English sentences: 'The door is green', 'The green door is locked', 'Green is a soothing colour' ... (Strawson 1974: 21) The argument here is straightforward: we need concepts, not merely to be things specified by concept-words in predicative position, but also to help us out when we have concept-specification in subject or adjective position. However, this is impossible so long as concepts are understood as essentially predicative or unsaturated, since as such they are fit only to serve as the referents of concept-words in predicate position.

In essence, this is a kind of argument that Wiggins uses, and it lies behind his stipulation that concept-words such as 'man' should be counted as referring to entities such as *man*, which are neither saturated nor unsaturated. Of course, Strawson is rejecting Frege's scheme, not trying to repair it. Frege grounds the predicative ingredient of the basic combination on the essential unsaturatedness of certain entities—concepts. Strawson instead speaks of our having to reflect at the level of concept-words the relation of exemplification that obtains between concepts as he conceives of them and particulars. Still, one might reasonably put this question to Strawson: when we do use a concept-specifying word—in whatever grammatical form it occurs—are we not thereby specifying some entity? Here the answer would seem to be affirmative: as earlier noted, Strawson has no hesitation in speaking of concepts as entities specified by concept-words like 'man'. These entities are not now to be understood as unsaturated entities, but they are entities nonetheless. And now a certain mystery emerges.

Since Strawson regards particulars and concepts as entities, and also insists, still at the level of ontology, that particulars exemplify concepts, we might well wonder whether exemplification by particulars is an essential or merely added feature of concepts. At the level of the sentence, there is little doubt that the feature which reflects or shows exemplification is something Strawson finds himself able to 'ascribe' to the predicate phrase—the phrase consisting of the concept-word and either the copula or a verbal ending (Strawson 1974: 30). Yet, how can being exemplified by particulars be a feature that just happens to follow our manœuvres at the level of expressions? Taking seriously the idea that concepts are entities, it surely must be the case that, if they have the feature of being exemplified by particulars, they have it intrinsically. (Of course, if exemplification were itself an independent

item, something that could be specified by its own concept-word, then we might imagine it being sometimes co-present with concepts like *man* and sometimes not. But, perfectly reasonably, Strawson resists this as leading to regress.⁴¹)

However, if exemplification is an essential feature of concepts—an intrinsic feature of these entities—then the Strawsonian picture begins to merge with Frege's. There is after all little to choose between concepts as unsaturated entities standing in need of particulars, and concepts as entities that are necessarily exemplified by particulars. But this is surely a consequence Strawson would have resisted, and it is easy enough to find the ground of his resistance.

The feature of Strawson's account that makes for trouble is his readiness, on the one hand, to allow concept-words to specify entities, and, on the other, his conviction that this ontological talk should not be taken all that seriously when we come to discharge the function of propositional combination. This mirrors Wiggins's allowing that concept-words like 'man' refer to the concept-entity *man*, while at the same time insisting that it is items in the semantics like 'is a man' and not concepts themselves that are unsaturated. However, unlike Wiggins, Strawson offers an explicit reason for his reluctance to take the ontology of concepts seriously.

He begins by imagining an objection to his speaking both about 'expressions specifying particulars' and 'expressions specifying concepts'. The worry behind this objection is that the second of these invites us to think there is ontological parity as between particulars and concepts: We return, then, to the complaint that the notion of concept is obscure in a way in which that of spatio-temporal particular is not. The complaint may take a more specific form. Spatio-temporal particulars, it seems, are not the creatures of language; they have their own independent being in the world, owing nothing to words (except when they are utterances or inscriptions of words). But there is no hope of understanding what concepts are except by seeing them as the creatures of language. They owe their being, such as it is, to words. So how could one hope to explain, in any genuine and non-circular way, the functioning of expressions by reference to the nature of concepts?⁴² To which he replies: 'The contrast is a good deal overdrawn. But let it stand. Let it be allowed that to talk of concepts is to talk of the senses of expressions' (Strawson 1974: 39). From this perspective, one can see why Strawson would think my argument simply missed the point. Necessary to that argument was our taking seriously the idea of concepts, and of their being exemplified by particulars, as serious

⁴¹ There is another possibility. Perhaps we should think that there is some third ingredient at the level of ontology—something which itself is not a concept, is separable from concepts and particulars, yet which somehow cements the one to the other. Some such proposal can be found in, e.g. Armstrong (1997: 114–16) when he speaks approvingly of the 'ontological copula'. I have not gone into this kind of possibility here because it is so alien to Strawson's outlook, and seems besides of doubtful coherence.

⁴² Strawson 1974: 39. You would not find anything like agreement about this rather rhetorical paragraph. For example, Armstrong (1997) insists that concepts—he prefers to speak of universals—so far from being creatures of language, are more solidly grounded than anything we could derive directly from language. For him, genuine concepts are known to us through science and not merely as by-products of our use of predicate expressions. Nor is he alone in maintaining some such line.

ontological matters. But here we find Strawson insisting that, for all he has allowed talk about concepts as entities specified by concept-words, this is in the end just talk. Hence, it is no wonder my concern about whether exemplification by particulars was an essential feature of concepts is not shared by Strawson: there is, in his view, no language-independent way to get at concepts, and certainly not at their putatively essential properties. The bottom line is that the functioning of predicate expressions cannot be revealed by the nature of concepts in the way that, one surmises, the functioning of subject-terms is revealed by the nature of spatio-temporal particulars.

It would be wrong to presume that Strawson's attitude toward concepts—thinking them merely ‘creatures of language’—is shared by Wiggins.⁴³ But, if it were, one could see how it would justify his insouciance in moving from Frege's talk of unsaturated entities to talk of unsaturated expressions.

Aside from the issue of Fregean interpretation, does it matter that the asymmetry between concepts and objects is no longer an ontological one? I think it does. As I have up to now only claimed, I think there is an important insight about predication in Frege's *unrepaired* scheme—an insight which is lost when one does not take seriously the ontological asymmetry of object and concept in that scheme. With Strawson's and Wiggins's views as background, it is now time to flesh out that claim.

2.4. Fregean Asymmetry

Why did Frege insist on there being an ontological basis to the asymmetry of concept and object? Or, since one could never be sure of Frege's intentions here, why *would* someone insist on this? There seems to be a pretty standard answer: in distinguishing saturated objects from unsaturated concepts, and then showing how they are made for one another, one has a way of dealing with the basic combination without raising questions about inertness or regress. The basic combination is the most primitive device for expressing that which is potentially true or false. If we allow the pairing of saturated object and unsaturated concept to constitute what is expressed by the basic combination, then there is no way for worries about inertness or regress to arise.

But this won't do as an answer to my question. I asked why someone would insist on *ontological* asymmetry, and not merely on asymmetry. In the basic combination one finds both a device for referring to some item and a device for predicating something of that very item. This is, if anything is, a truth universally acknowledged. But, as is evident in the Strawsonian account, and also in the Fregean scheme after Wiggins's repair, these needs can be met by distinguishing asymmetric kinds of

⁴³ It is very important in understanding the ‘creatures of language’ claim to keep in focus the fact that Strawson is talking of his (and Wiggins's) notion of a concept and not Frege's. On some readings of Frege's notion, a concept is very much like, perhaps even the same as, a property. Yet, though there are issues about the ontological status of properties, few would assent to the claim that they were simply creatures of language and owed their existence to words. That said, it is surprising that Strawson doesn't so much as mention the notion of a property, either in this context or anywhere else in his book.

expression: a subject-term to accomplish the necessary reference and a predicate expression which completes the job. When this is all laid out properly, the basic combination can be seen as expressive of the true-or-false, and the problem of regress need never arise. So, simply coping with the basic combination by positing some kind of asymmetry does not explain why anyone would insist on that asymmetry being ontological.

There is an obvious supplement to the previous answer, and it might be thought that this does the trick. Notoriously, Frege was in the thrall of a certain model of language use: the name-object model. He seems to have felt that, whatever the linguistic function, it must be explicable by something along the model of a name's capacity to refer to a particular. This is just fine when it comes to the referential functions of words, but in respect of their predicative function, the model is put under strain. Attempting to cope, the story goes, Frege found himself forced to discover an entity that is at once the referent of a predicative expression while also capturing its predicative function. The result then is the problematic notion of an unsaturated entity.

There is something in this thought, but I still do not believe it gets us to the bottom of Frege's insistence on the ontological asymmetry of object and concept. It is certainly true that Frege extended the name-object model to what many regard as its breaking point, that is, its application to whole sentences, and even its less extreme application to predicates has kept Frege's interpreters and apologists busy. But there is still something unsatisfying about this as a complete answer to my original question; something more seems to be at issue.

Though certainly not his way of putting it, I believe, first, that Frege saw there to be a crucial and radical difference between *what is accomplished* by subject-terms and *what is accomplished* by predicate-terms, and second that, while radically different, these two accomplishments must be treated in the same way. Difference is of course revealed in the differing linguistic apparatuses we deploy in when reference and predication are expressed in the basic combination. That is blindingly obvious. But Frege also saw a need to base the difference on something independent of these apparatuses—on something, that is, which is not essentially linguistic. This was because he sought a way of treating reference and predication as separate but *equal* contributors to the basic combination, and one of them, reference, had an obvious basis starting from outside language, that is, the name-object model. Since the name-object model was the tool he envisaged for describing, from the outside, how language works, his options for providing the same kind of ground for predication were limited. In order to explain—extra-linguistically—the radical difference between what is accomplished by subject-terms and what is accomplished by predicate-terms, while maintaining parity of treatment, he had to find very different entities corresponding to each. But, as we have seen, any such manoeuvre better suits reference than it does predication. So, although the aim was parity of treatment, the result was disappointingly one-sided and, as Wiggins and Strawson both recognize, does not ultimately cohere with his overall project.

Still, the crucial insight that I find in Frege's scheme survives the havoc wrought by the name-object model. The basic combination brings together two fundamentally different and complementary tasks: reference and predication. It seems therefore only right and proper that any account we give of these tasks will treat them with equal seriousness, even while showing them to be radically different. Frege's distinction between two fundamentally different kinds of entity—objects and concepts—can be understood as aiming to achieve just that. In contrast, proposals like that of Strawson and Wiggins simply do not take the 'separate but equal' requirement on board; separateness is catered for but not equality, and this lack of equal treatment is not even noticed.

The referential function of the basic combination is explained from a starting point outside our use of subject expressions. Reference is seen as providing a special kind of link between these expressions and spatio-temporal particulars, and because particulars do not owe their existence to our use of such expressions, the link itself is not essentially tied to language. However, though lip-service is paid to an extra-linguistic starting point in respect of predication—I have in mind here Strawson's notion of a concept and its exemplification by particulars—in the end predication comes to be understood as a task internal to our use of certain expressions. Predication is what you have when a concept-word is combined with either the copula or some appropriate verb ending, and the concepts called on here by concept-words, so far from being independent of the employment of expressions, are characterized as mere 'creatures of language'.

Given what was documented at the opening of section 2.3, namely, the near universal recognition of the complementary roles of reference and predication in the expression of truths, this oversight should be surprising. The differential treatment of reference and predication should, one thinks, at least be noticed. But of course it isn't, as witness the effort I have been expending to so much as make the subject of that oversight visible. Nor am I confident of having yet expended enough effort in this regard. So, leaving Frege and the others on one side for the time being, let me approach the matter from a different angle. My aim will be to engage in a bit of consciousness raising about the unequal treatment given to predication.

2.5. Fair Play for Predication

I begin with some straightforward (i.e. rough and ready) observations about reference. When someone is told:(Ref) X refers to O, what would you predict the hearer would take X to be? A person, a word? Something else? The most likely answer, outside of philosophical discussions, would be the first, that X is taken to be a person, though of course all three are possible. Philosophers find it more natural than ordinary speakers to think of X as a word or phrase, and, though less common, there is nothing unlikely about cases in which X is neither a person, nor a word, but some object. The bundles of everyday things

used by the Academicians of Lagado (in Swift 1726/1985⁴⁴), the little models in the Paris courtroom appealed to in the reconstruction of an automotive accident (which, as already noted, captivated Wittgenstein) and, more mundanely, the perfectly ordinary use of items on the dinner table to tell a story—these are just some of the contexts in which it might be said that X, a non-word object, refers to some other object O. Much less intuitive, though in philosophy almost anything is possible, is the case in which X is taken to be a constituent of thought, perhaps, as has been maintained, some token in Mentalese. Though it is not difficult to imagine substantial reservations about this understanding of X, the fact that one comes across it serves to show just how promiscuous (Ref) is. For, though Mentalese is held to be a sort of language, tokens in it are certainly not words in any ordinary sense.

What this range of cases shows is something even stronger than was suggested in the previous section. It is not merely that the task of reference can be explained from a starting point outside language—that is, one beginning with the independent existence of spatio-temporal particulars—but it is even possible to think of reference as functioning in complete independence from words. Nor is such non-linguistic reference in any way parasitic upon its achievement by linguistic expressions.

Can these same things be said of predication? That is, can we understand predication as a task explicable either from a starting point outside language or even, more radically, independently of language? In so far as reference and predication are accorded equality of treatment, the answers here should certainly be affirmative. But there are obstacles even to seeing how to put these questions in the rough-and-ready way that served for reference.

The immediate problem is terminological. The points made about reference began from claims summarized in (Ref), but there is no similarly neutral way to summarize claims about predication. Thus, leaving aside its awkwardness, we would expect someone hearing: X predicates of O, to think of X as a word or words. Slightly better would be: X describes O, where we can perfectly well think of X as a person, but even here words are not far away. For when we think of how X would have set about the task, we undoubtedly think of the use of a word or words.

The simple fact is that there is no term in general employment that we can use to characterize the predicational task that does not put us in mind of expressions in a language. Quine said, ‘predicates are the other parties in predication’, but, in every primer of philosophical logic, a predicate is, or is constructed out of, words in natural language. Wiggins, perhaps more helpfully than Quine, says that ‘predicates

⁴⁴ Later, and at the risk of ruining a good story, I shall look more closely at *autonomia* or self-naming that is so rampant in the Academy Swift describes. A satire on philosophers it may be, but Swift's story has more than one important philosophical insight.

describe', but this is little better for our purposes, since we are bound to think of a *description* as something undertaken with words.

It might be thought revelatory of something deep about predication that we have a problem finding a label for it which doesn't carry the suggestion that words are involved in its accomplishment. Perhaps this is even further reason to think, with Strawson, that concepts are 'creatures of language'. Well, no. In fact, I think the truth is rather quite the reverse: the problem of terminology doesn't *reveal* something deep about predication, but rather *reflects* a widespread and systematic misunderstanding of it. In the next section, I shall offer a suggestion for supplementing our terminology, showing thereby a way to liberate our understanding of predication. But, as part of preparing the ground for this, I shall first probe this misunderstanding a bit further.

Frege's insight about predication was genuine, but it came out all wrong, perhaps because he was in the thrall of a name-object model of language. Aiming to show that predication is radically different from, but on a par with, reference, Frege's use of what is a referential model of language-use was bound to lead to trouble. It left him only one way in which to capture the difference, namely by finding the objects singled out by referring expressions to be radically different from those singled out by predicates. Most certainly, Strawson and Wiggins are not in thrall of the referential model. But there is still a sense in which they give the notion of reference more than its due, and they have this in common with the whole community of philosophical logicians.

Reference and its correlative ontology are taken seriously; the objects of reference—spatio-temporal particulars—are thought of as language-independent entities, and reference tends to be treated as foundational for our view of word–world relations. In contrast, predication, and the objects it calls on, are not accorded anything like the same linguistic or ontological status and this leads to a kind of inequality. When the standard of language independence is set by reference, predication itself becomes a 'creature of language'. All of this is shown in schematic form in Figure 2.1. Begin with the horizontal line which divides the level of language (L) from a level that, with some trepidation, I have labelled a level of ontology (O). Above the line is the basic combination, the *linguistic* structure which expresses what Strawson calls our 'primitive thought about the world'. Below the line is a structure capturing the subject matter of that thought. (My trepidation about the label 'ontology' comes precisely here: Strawson always speaks about our 'thought' about the world, and not simply about the world. But just for the minute, there is no need to worry too much about the label. Think of it as basically a level of *language independence* rather than as requiring heavy-duty metaphysics.⁴⁵ Note too that I have put 'exemplify' and 'concepts' in parentheses, suggesting thereby that their ontological status might well turn out less robust than that of particulars.)

⁴⁵ If my interest in Frege was scholarly, this would be the place to talk about the ontology of semantic values and what Frege means by 'Bedeutung'. But nothing I go on to say requires me to engage that deeply with Frege's texts.

of predication; it is there to smooth over uses of concept-words when they occur in contexts other than of the predicate expression, for example, when the concept-word occurs in subject position. There is in fact no representation of the task of predication in Figure 2.1, and this is at the root of my complaint. Reference and predication in Figure 2.1 are treated differently—and this is of course right—but they are not treated equally. Reference, as we have seen, is treated as a language-independent task which happens, in respect of the basic combination, to serve a linguistic purpose. But the representation in Figure 2.1 makes it impossible to conceive of predication as a task that could, in this way, be independent of language.

A Fregean picture in which there is a referential relation between the whole of the predicate expression and a very special sort of entity could easily be superimposed on Figure 2.1. This would constitute at least an attempt to treat reference and predication equally. But, while I think Frege's instincts were right, adopting his ontological move here would be a step backwards. What we need is some other way to display reference and predication as complementary but equal contributors to the work of the basic combination. The alternative that I suggest is shown in Figure 2.2. In comparison with Figure 2.1, the axis of this scheme is rotated. Levels (L) and (O) are now side by side and the horizontal line has a different purpose: it marks off two functions or tasks. The first of these is of course reference, thinly disguised as 'Ref'. (A reason will be given for this disguise.) As before, the arrow schematically indicates an appropriate linkage between certain elements at level (L) and entities at level (O). But more important than these superficial differences is what happens below the horizontal line. Here 'Pred' labels the task accomplished by full-blooded predicate expressions, and, while this task is schematically shown by the arrow connecting these expressions with entities at level (O), the arrow used is not the same as the one used for Ref. This difference is intended to be more than arrow-deep. Whereas the Ref linkage is one that takes us from an expression to an object, thereby, as it were, *bringing an object to our attention*, the Pred linkage has a very different and complementary function. Instead of singling out an object, the Pred task consists in *information contained in the expression at level (L) being brought to bear on*

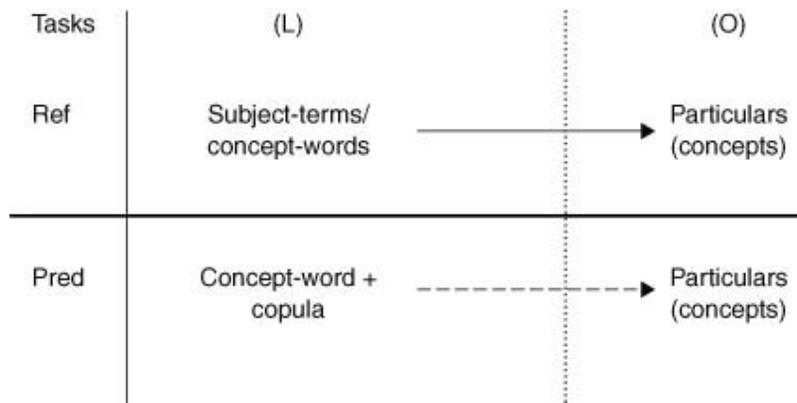


Figure 2.2

some relevant object. The basic combination can then be seen as the result of putting together the very different Ref and the Pred linkages. Thus, the expression ‘Socrates’ takes us to a particular person, and the expression ‘is bald’ brings information to bear on this very same particular. When Ref and Pred are co-ordinated in this way, the result is the structurally primitive thought that Socrates is bald.

Clearly, there is a lot more to be said about the Pred task, but not quite yet. What I should like to do first is to outline several substantial reasons for regarding Figure 2.2 as an improvement on Figure 2.1. Top of the list is the fact that predication is portrayed as fully on a par with reference. I used the label ‘Pred’ to avoid the suggestion that comes with ‘predication’, namely that it is an essentially word-involving task. As displayed ‘Pred’ is a task that, while it happens to be fulfilled by certain expressions (full-blooded predicate terms), could nonetheless be fulfilled, like reference, without calling on words. We have wordless reference when someone uses a salt cellar on the dining table in telling a story, and what Figure 2.2 allows is that Pred similarly might be fulfilled wordlessly. (I would have said that Figure 2.2 allows a place for ‘wordless predicates’, but this just sounds incoherent. That is why I have had to ‘disguise’ predication as ‘Pred’ and, having seen the need to do so, why I had in fairness to disguise reference, even though there is no whiff of incoherence in the expression ‘wordless reference’.⁴⁶)

Consider next an obvious difference between the figures that has not so far been mentioned. Figure 2.1 contains a structure—something at the level of ontology (O)—which grounds the basic combination. In effect, the basic combination is treated as a device at level (L) that *expresses* the level (O) exemplification of a concept by a particular. But there is nothing in Figure 2.2 at the level (O) which matches this; there being at that level only spatio-temporal particulars and, perhaps, concepts. In Figure 2.2, what plays the grounding role for the basic combination is the juxtaposition of functions or tasks fulfilled by relevant subject and predicate terms. The basic combination is a basic (or primitive) combination of two different and equal semantic tasks. Now it might be thought that, while there are some advantages to

⁴⁶ Here is as good a place to comment on a reservation that someone might have about the whole idea of wordless reference. Consider the example of the dinner party story in which a salt cellar is used to refer to a car. Isn't it the case that this example only works because, in telling the story, one says, ‘let this be my car’ (or words to that effect)? Doesn't this show that the referring capacity of objects is in some way parasitic on the referring capacity of words? No, it doesn't. Admittedly, in the case described, words help us fix the referent of the object, but this is no surprise and certainly doesn't take away from the fact that the capacity to refer resides in the object. On the one hand, it is certainly possible to imagine that the reference of an object is established without recourse to words; think here of how we might use gestures or even simply of certain salient juxtapositions to assign or comprehend the referents of objects. On the other hand, it should come as no surprise that, in the case imagined, we use words to guide our audience, rather than simply leaving them to work out for themselves what is happening. It is after all a *story* that is being recounted in the dinner table example. (While I do think that there can be wordless stories, and even think that these are important to us, most of the stories we ‘tell’ are worded.) Note finally that nothing I say would be undermined by its being true, as it probably is, that only language-using creatures are capable of using objects as referring devices. Clearly, I intend the ability to refer, whether with or without words, to be a semantic ability, and as such it is probably only available to semantic creatures like us. But this is not to say that any individual act of reference using a non-word object is parasitic on linguistic reference.

treating the basic combination in this way, there is still a loss: we now have no structure that counts as being expressed by the basic combination. However, what appears to be a loss is actually a substantial reason for preferring Figure 2.2. The idea that the basic combination expresses some structure at level (O) has never been taken all that seriously by those who share Strawson's reservations about the ontological status of concepts and the relation of exemplification. Moreover, there is something suspect about the use of the level (O) structure in Figure 2.1—something only implicit in what I have said up to now—that makes the move to Figure 2.2 even more attractive.

Think back to Strawson's discussion of subject-terms and concept-words. Subject-terms *refer* to spatio-temporal particulars, whereas concept-words merely *specify* concepts. Nonetheless, specification still counts as a kind of referential relation, and it is this whiff of reference that is ultimately responsible for the somewhat embarrassing presence of concepts on the ontological side of the divide. Concept-words specify (i.e. sort of refer to) concepts, but the latter are, if Strawson is right, 'creatures of language' rather than independent existences. Even more embarrassing is the notion of exemplification. It seems to be a relation—a kind of concept—but we are told that we should not take this at face value. For if we try to find some expression which specifies this concept—perhaps, in the primitive case, the copula—then we will end up with a hopeless regress. We thus find ourselves *talking* about exemplification, but cannot take ourselves to be thereby *specifying* the concept of exemplification. Still, the embarrassment caused by concepts and exemplification is easy enough to overlook, because what seems important are not these ingredients so much as the structure in which they figure: the exemplification of a given concept by a given particular. It is this latter structure that is expressed by, and thus grounds, the basic combination. The idea that the basic combination expresses some item—perhaps a possible state of affairs—is familiar and it seems therefore easy enough to accept, leaving any wrangling about metaphysics for later. However, in the present context, I think we should be just as worried about the idea that a sentence *expresses* a state of affairs as we are about the idea that a concept-term *specifies* a concept. Nor is the problem metaphysical.

Frege perhaps overextended the referential model of name and object: he thought we had to provide reference for concept-words and sentences just as we do for names. As noted, he also thought we could accommodate the reference–predication distinction within this scheme by being careful about the kinds of thing we allowed in our ontology. But the result of these manœuvres turns out to be incoherent, and not merely ontologically suspect. Strawson and Wiggins are most certainly *not* in the thrall of name-object model, but there is a sense in which they still allow reference to play too large a part in the enterprise.⁴⁷ For just as there is more than a whiff

⁴⁷ Perhaps I have left it a bit late to say this, but we of course shouldn't simply assume that the scheme Wiggins comes up with is one he himself would endorse. He may, and if he does, he is subject to the same worry I have about Strawson's view, namely that it gives reference too large a role. But Wiggins was after all only setting out to repair Frege's scheme—to make whatever minimal changes are necessary to render it coherent—and thus the central role of reference is unsurprising.

of reference in talk of concept-specification, so there is in the idea of a sentence expressing some state of affairs.

Looked at carefully, the whole of the scheme shown in Figure 2.1 is shot through with the idea that the key relation between words and the world is broadly referential; names refer, concept-words specify, and sentences express. In a word, this scheme has a built-in bias to reference. Hence, it should come as no surprise that we find it difficult to place predication in that scheme in a way that displays equality of status with reference. It is thus no accident that Figure 2.2 has no place for special structures such as states of affairs that are expressed by sentences. The problem is not that these structures are ontologically suspect; it is simply that they make impossible parity of treatment as between reference and predication.⁴⁸

As has been said more than once, what is presented in Figure 2.2 is schematic. Depicting Ref by a solid arrow and Pred by a dotted one *displays* them as different, but does nothing to spell out what this difference is. My one attempt to do that—my claiming that Ref linkages bring particulars to our attention, while Pred linkages bring information to bear on particulars—needs (and will be given) further elaboration. Still, even with this unfinished business, it is possible to see Figure 2.2 as an improvement on Figure 2.1. For, to repeat, it displays the ingredients of the basic combination as distinct, complementary, and, most significantly, as equal. Ref, the thinly disguised task of reference, is brought to bear on particulars, and so is Pred (though there is more reason for terminological disguise here). When we have an appropriately co-ordinated exercise of these two tasks, we have the basic combination. Figure 2.2 thus captures the idea that names and predicates ‘are made for one another’, without the distortion that comes from locating predication in a world of reference.

In displaying the parity between Ref and and Pred, Figure 2.2 makes room for the possibility that items other than words can fulfil both of these tasks. However, while there are clear examples of particulars being drafted in as referring devices, I have not so far shown that this is more than a possibility in respect of Pred. Adding this to the issue of terminology and the need to say more about the task of Pred, the list of unfinished items of business is now substantial. Still, with these questions about the Pred task hanging in the air, we have in a real sense rejoined the main theme of this chapter.

I began by asking whether objects might, in appropriate circumstances, take on the functions of words. It was clear at the outset that this is unproblematic in the case of reference; predication has been more of a challenge. I opened the discussion with Goodman's notion of exemplification since it suggests the possibility of objects helping out predicates. However, before I could take this further, I had to say more about predication itself. This is because, while lip-service is played to its independent role in the basic combination, predication is usually accommodated within a

⁴⁸ This is intended to be concessive, but only in this context. I actually think that a lot of contemporary metaphysics with its talk of universals and ontological copulas is really little more than the result of the shadows cast by a bias towards reference that passes for analysis amongst a certain community of metaphysicians.

framework that takes reference as in some sense central.⁴⁹ Now that we are in a position to conceive of predication as it ought to be—as a task on a par with, though radically different from, reference—we are just about ready to return to the central question of the chapter, namely whether an object can take on the function of a predicate. But I will rejoin this question by first introducing some terminology. This is to prevent us being distracted by the word-involving resonance of ‘predicate’. Also, I will have to say more about our understanding of the Pred task.

2.6. Predication by Another Name

In school grammar, one learnt to say that adjectives qualify nouns. Presumably, the idea is that an adjective adds some quality or qualification—something further or more precise—to whatever the noun introduces. Leaving on one side the adequacy of this as a grammatical truth, it will be convenient here to borrow the expression ‘qualification’ as the label required for my purposes. There are two good reasons for this: first, while close enough to predication for the connection to be intelligible, in being old-fashioned, specialized, and generally out of use, the word itself does not suggest any of the word-involving prejudices of ‘predication’. Second, leaving behind its grammatical origin, the label has resonances that are extremely useful in the present context. In particular, it allows something like the same latitude in use that ‘reference’ does. Thus, given the claim: *X* qualifies *O*, one might easily and naturally think of *X* as a person or word, and, as my examples shall show, one can as well think of *X* as an object. To be sure, when we speak of a person *X*, or object *X*, qualifying an object *O*, there is bound to be uncertainty about exactly what is being said. I will of course address this concern below. However, the point here is simply that the locution is not odd. (The contrast here is with the distinctly odd: ‘Person or object *X* predicates *O*’.)

⁴⁹ I emphasize again that my discussion of Frege, Wiggins, and Strawson should not be thought as an exhaustive account of these matters. I focused on the issues raised by their work because it seemed the quickest and most perspicuous way to make my point about predication. However, there are other ways of looking at these matters that I never touched on. One such way is Wittgenstein's Tractarian idea of treating predicates (specifically, relations) as items shown by the arrangement of objects in a proposition. However, any discussion of the *Tractatus* would, in the context of this book, be wholly superficial, so I haven't attempted it. In any case, I really don't know what a Tractarian object is—though I am fairly sure it is not what I mean by this term—and I am not alone here. Of less purely historical interest is the treatment of the basic combination by means of differential clauses in some Tarski-style theory of truth. In the most primitive, non-quantificational case, reference is handled by one kind of clause and predication by another. This gives the appearance of equal treatment to the notion of a predicate, but there are problems here that turn on how we understand the notion of satisfaction that figures in the clauses governing predicates. On one way of understanding what is going on, we have what is in fact a version of my Figure 2.2. But I do not think that this is how the difference between reference and satisfaction clauses tends to be viewed, and a deeper look at these issues would take me too far away from the business of this chapter.

Alright, so we have a label and I have sketched what the label is intended to do: it marks a task almost universally thought of as something done with words in specific constructions in natural language, but which, under this label, should be thought of as something that could be accomplished without words. But what exactly is this task? And when we know what it is, are there any interesting cases in which persons or objects (and not merely words) can be said to qualify something?

Deferring the provision of examples to the next two sections, let me do my best here to say more about the task of qualification. In the previous section, I characterized predication as the bringing to bear of information on particulars. In the special case of a fully linguistic predicate such as ‘is a man’, we look to what is generally called a ‘theory of meaning’ for a more specific characterization of the relevant information. Yet the point can be made in advance of settling on any specific theory, or even settling whether such theories are a good idea. Perhaps through mastery of conventions, truth conditions, or perhaps in some radically other way, speakers of English can be described as having the capacity to bring to bear the information associated with a predicate on relevant spatio-temporal particulars. In a sentence such as ‘Socrates is a man’, our mastery of ‘is a man’ makes available information that happens to be brought to bear on Socrates (via ‘Socrates’), but this same information could have been brought to bear on a whole range of other particulars.

This characterization of predication carries over to qualification, though of course, in making this transition, we have to leave theories of meaning on one side. Thus, we can say that when an object X qualifies O, X either brings, or is intended to bring, information to bear on O. Since X might be a non-word object, we must be prepared to tell a story about the nature of the information associated with X and brought to bear on O which is substantially different from the one told about linguistic predicates. Some idea of how this might go will be clearer with the examples I shall offer in the remainder of this chapter and the next one. However, the important point here is the recognition that the *task* of predication is one and the same as the *task* of qualification. As already noted, the model here is reference. Reference is the same kind of activity, whether it is achieved by means of words, objects, or elements of thoughts. ‘Reference’ is thus a superordinate category—a general name of a task—under which we can group the systematically different ways in which this task is carried out. In exactly parallel fashion, I intend ‘qualification’ to be the label of the other task in the basic combination—a task which, in that combination, is unsurprisingly attempted by words. In effect, ‘qualification’ is superordinate, and ‘predication’ labels that same task—the bringing to bear of information—by means of words in natural language.

Being so unfamiliar, my account of qualification is liable to be misunderstood. However, the following notes should help:

- (i) Talk of the information carried by objects might all too readily put one in mind of rather technical ideas about information theory and/or familiar stories about tree rings and rain clouds. Yet it is important to see, even before examples are discussed, that these play no part in my understanding of

qualification. It is true enough that the number of tree rings informs us about the age of the tree and that the presence of certain kinds of clouds inform us about the likelihood of rain. But in neither case is there anything like a predicative relationship between the informing item and the item about which we come to be informed. One way to put the difference is this: the tree rings do not bring information *to bear* on the tree; information about the tree is *extracted* from them.⁵⁰

- (ii) I can imagine someone complaining that talk of 'bringing to bear of information' is unhelpfully close to the idea of predication, and is therefore of little explanatory value. However, while I think there is something in this complaint, I don't think it damaging. Think about the ways in which we tend to characterize reference. We say that N refers to O when N picks out O, or when N labels O, or when N stands for O. Each of these is perilously close to the original notion of reference, but some kind of circularity here seems unavoidable. Attempts to say what reference is in completely other terms tend to lose track of the thing itself. (I cannot of course argue for this here, but offer as some evidence the fact that, in spite of the effort expended, there is simply no extant proposal that is even remotely plausible. Straightforward causal accounts just don't work, and appealing to speakers' intentions, as one is forced to do, reimports reference, albeit at the level of thought.) My suggestion then is that this same rather profound circularity infects attempts to say what qualification accomplishes. We can say that X qualifies O when X brings information to bear on O, when X describes O, when X characterizes O, when X is true of O, and so on. None of these would suffice to explain what is going on to a creature who had never encountered the notion in the first place. But this is just how it is with both reference and qualification.
- (iii) My insistence on treating qualification as *different* from but *equal* to reference should not be mistaken for treating them as *independent* of one another. I do think that human beings have in their repertoires two semantic abilities: the ability to use objects or words-objects to refer to other objects; and the ability to see in objects or words-objects a potential for informativeness, an aptness to serve as sources of information that can be brought to bear on other objects. I also think that the second of these abilities has not been given its due, largely because, when it is exercised in natural language, it tends to be spelt out in terms of reference (often trading under the label 'concept specification'). But I do not think that we can exercise these two abilities independently of one another, or independently of the truth-directed basic combination. Indeed, I would argue that a creature only has the capacity to engage in acts of

⁵⁰ It might be tempting to think that the distinction matches Grice's between natural and non-natural meaning. But while there are connections here, the two distinctions are not the same. Grice's distinction is essentially that between something we do and something we find. Clearly, predication and qualification belong with the former, but, in so far as qualification allows non-word objects to figure in our actions, what we find in them—or perhaps even put into them—is crucial.

reference if it also has the capacity to engage in acts of qualification, and vice versa. We cannot discern the one ability without the other. Moreover, each joint exercise of these different abilities *is* the production of a truth-directed structure, whether in words or thought. It is not that we just have two semantic abilities which, rather like the result of a chemist mixing substances, happen to produce some third thing, something apt for being true or false.⁵¹

There is more to be said about qualification, but it is best said in the context of actual examples. In any case, unless there are interesting cases in which certain objects can be said to qualify others, the notion of qualification would be little more than a curious possibility. (It will prove anything but a curiosity, so the label ‘qualification’, whether it resonates or not, will bulk large in what follows.)

2.7. Initial Examples

The initial pair of examples will seem familiar.

1. You are in a city in the Far East where your language is not spoken. Passing a shop whose window displays all manner of men's suits, the proprietor gestures for you to stop. He is holding a book of swatches of cloth that he has opened to a particular place, and he excitedly points at the swatch on that page, while looking back towards his shop windows and entrance. I say that in this case the swatch qualifies a suit he proposes to make. (One could also say that the proprietor qualifies a suit by using the swatch, but the focus here will be on objects as qualifiers.⁵²)
2. You receive a parcel of information from an estate agent about a flat you are thinking of renting. In amongst this information, you find a single sheet of paper on which are mounted small square coloured cards. There are captions under each card, for example, a caption under one reads: ‘bedroom 2’. I shall say in this last case that the coloured card qualifies that bedroom.

It should be said at the outset that these examples are problematic: being clearly adapted from examples that Goodman uses they are bound to make one wonder whether, in spite of the build-up, qualification is simply exemplification by another

⁵¹ This book is not the place to argue for these interdependencies, but by asserting them I hope to defuse irrelevant objections to the notion of qualification. No one doubts but that reference and truth are intimately linked in the basic combination, and even if you regard truth as somehow basic (see Davidson 1984*b*), reference doesn't simply disappear. My suggestion is simply that we widen the circle of intimacy a bit so as to include qualification.

⁵² Pretty clearly, any case in which a person qualifies an object will be one in which a person uses some object or prop to do so. This in no way ruins the parallel with reference: in any case in which a person refers to an object, I think you will find that there is some object or prop (perhaps a word, perhaps a gesture) by which the reference is effected. There are issues for both qualification and reference when one tries to imagine cases in which these tasks are undertaken in thought, but they are not relevant here.

name. Nor is this the only difficulty. When used as examples of objects serving predicational roles, they have two related weaknesses. First, both seem to work only because they are set in highly conventionalized contexts; this suggests that the phenomenon of qualification is unlikely to be general enough to be interesting. Second, they seem to depend on natural language predication in a way that might undermine their claim as examples of qualification—examples in which the predicative function is discharged by non-word objects. While admitting that these are not the best examples of qualification—better ones will follow in the next section—I should like nonetheless to address these difficulties, not least because it will allow me to reconnect with my earlier, inconclusive, discussion of Goodman's notion of exemplification. Let me begin with some comments about the role of context in these examples.

Context is going to count for a lot in specific examples of qualification, but that fact alone shouldn't count against those examples; some kinds of context dependency are perfectly harmless.⁵³ Just to take the first example: unless you knew about clothing, tailors, and perhaps even about the bespoke tailoring industry that exists in certain countries in the Far East, the scenario you witnessed would strike you as simply bizarre. Yet I doubt that any worries we might have about qualification in these examples is based on the need for some such general social setting. After all, it is widely accepted that the same need exists even for predicates that are unproblematically linguistic.

There is a second strand of context, perhaps even more crucial to the examples, and even though there is no similar appeal to context in linguistic predication, the dependency in cases of qualification is harmless. Think of what we would have to know (and do know) in order to recognize that the informational target of the swatch is a single item—a suit—and that, in effect, the swatch, in so far as it is a predicate, is a monadic one. Clearly, we get information like this pretty much for free (i.e. non-contextually) in linguistic predicates: there are one, two, or more places or 'slots' which we recognize and which tell us the predicate is monadic or dyadic, etc. Objects, however, don't have slots, and we therefore must depend on context to tell us whether information in them is brought to bear on single items or pairs, etc. (The examples of object-qualifiers given in this chapter will be monadic, but there is no deep reason for this. A wider range of examples will be considered later.)

It is a third strand to the notion of context which is I think responsible for the worries one might have about the two examples. Not only do we need to understand something of the social background, not only do we have to look to context to fix the predicational domain, we also have to understand something of the conventions that govern books of swatches, colour cards, and other similar devices. In particular, we must understand that the objects which figure in the examples come from series of similar objects, and these series are conventionally used to provide a

⁵³ Much more will be said about context when we come to consider metaphor itself. So the comments which follow here are only a start.

certain kind of information—a kind that is often linguistically specifiable. Thus, suits are made of a certain fabrics, and the sample book is the conventional way in which we come to understand which *fabric*. Similarly, all manner of objects are coloured, and colour cards or charts help us pinpoint precisely which *colour* is in question. Given this, it would be natural enough to think, on the one hand, that qualification as illustrated by these examples is at best a highly restricted phenomenon; and on the other that any such qualification is parasitic on linguistic predication. With the second of these we in effect return to the issues raised by Goodman's notion of exemplification. But the worry about the restrictedness of the examples should be addressed first.

If all examples of qualification were dependent on the conventions that govern the many versions of what I shall call 'sample series' cases—swatches, colour cards and charts, differently stained slices of wood, wallpaper books, etc.—then the phenomenon of qualification would be less interesting than I think it is. However, in the next section I will consider whole ranges of examples which in no way involve such series or such conventions, so I will let them make the case for the pervasiveness of qualification. Still, I should like to say something here by way of opening the account in favour of these admittedly restricted examples.

Sample series are governed by conventions about how we are to arrange and use relevant sample objects. But of course similar conventions also figure in respect of linguistic predicates. Consider what we have to learn, for example, to use the word-object 'is a man' in application to certain particulars. Competence with items in the lexicon require, among other things, mastery of conventions that are quite as specific as the conventions governing sample series. Given this, instead of thinking poorly of the initial examples, depending as they do on such specific conventions, one might think that the examples actually bring out the parallel between purely linguistic predication and qualification. The idea would be that the sample series conventions *mimic* the conventions that govern a typical lexicon. Unfortunately, this point cuts both ways. Someone might take the fact that the conventions in the sample-series cases parallel lexical conventions as leading us straight back to the other worry about the examples, namely, that they show qualification to be parasitically dependent on 'proper' predication. With a view to overcoming this worry, I now return to Goodman's notion of exemplification.

Goodman described the relationship between a predicate and an object which exemplifies it as doubly referential: the object must be in the predicate's extension—in this sense the predicate refers to it—and the object itself refers back to the predicate. Goodman never seems to have envisaged that there might be a wholly predicational, as opposed to a referential, role for the exemplifying object. Indeed, I suspect he would have thought that the sample-series cases might seem to work as predicates only because they depend referentially on linguistic predicates, and it is the latter that do the actual work of predication. Some evidence for this comes from remarks he makes in *Of Mind and Other Matters*. He writes: 'I like to keep the term "true" for statements. Statements in a language are true or they are false. I don't like to speak of a picture as being true or false, since it

doesn't literally make a statement' (Goodman 1984: 196). And in an earlier passage, he gives a reason for this: A picture like a predicate may denote certain events ... When the predicates in a text denote those same events ..., the picture and the text are to that extent inter-translatable; and the picture, though it makes no statement, might be derivatively called true or false according as the text is. But we must not forget that, strictly speaking, calling a picture true or false is false. (Goodman 1984: 98–9)

These passages concern a special class of objects—pictures—and it might therefore be felt that they are tangential to the issue of whether objects can serve as predicates. But of course as the second passage reminds us, a picture is a picture of something; Goodman regards depiction as yet another referential relation. Given this, one might well ask—and this is of course what I have been encouraging—whether an object *depicted* can serve a predicative function. These passages suggest that Goodman's answer would be unequivocally negative: when it comes to being true of something, only the predicates of a language will do. The swatch might well apply to the same objects as does the predicate, but it is to the predicate that we look for the contribution to truth, not to the swatch.

In section 2.2, I pointed out that it is often true that exemplifying objects help out with the work of the predicates they exemplify; *pace* things that Goodman suggests, exemplifying objects are, sometimes at least, not merely referential. However, I there said little about the nature of this help, making only the vague claim that exemplifying objects might offer additional epistemic routes to the information contained in linguistic predicates. I should like now to do better, not least because what I have to say should significantly increase the interest of my two examples of qualification.

As per an earlier example of Goodman's, suppose that someone asks about the colour of your house, and you answer this way: (C) My house is ... *here you hold up a colour card* ... this colour blue.⁵⁴ Goodman would say of (C) that the concept-phrase ('this colour blue') in the full predicate ('is this colour blue') refers to the card and, if true, also has the house in its extension. Indeed, it is partly for this reason that the card (and perhaps the house) exemplifies this concept-phrase.⁵⁵ Also, as the passages above suggest, he would insist that, in exemplifying the expression 'this colour blue', the card might well

⁵⁴ In what follows, I am not going to address directly the fact that this sentence uses a demonstrative. I shall have more to say about this rather special demonstrative construction in the next chapter, but it would only complicate matters to open that discussion here. In any case, one could imagine a slightly more complicated example to the same purpose which used a descriptive phrase in place of the demonstrative.

⁵⁵ Goodman clearly subscribes to some such picture of the ingredients of the basic combination as one finds in my earlier Figure 2.1. Moreover, he is not particularly careful to distinguish full-blooded predicates from the concept-words and phrases they contain, and he doesn't make much of the distinction between a predicate's being true of something and its referring to it. None of this matters for my discussion, but I have in this opening sentence tried to keep things tidy.

make manifest or concrete the contribution of that expression to the truth conditions of (C), but that it is simply a mistake to think that the card itself makes the kind of direct contribution of a predicate. My quarrel is with this last claim: I agree that the colour card concretizes and makes manifest something—that it can be said to enhance our understanding of the predicate expression—but I think the right way to look at this will bring us straight back to the un-Goodmanian idea that objects can themselves be predicates.

Imagine that instead of (C), you had answered: (C') My house is a light colour blue. What is the adjective 'light' doing here? Leaving on one side the thorny issues raised by certain sorts of adjectival construction, it seems obvious enough that 'light' makes a predicational contribution to the truth conditions of (C').⁵⁶ One way to put this would be to say that the expression 'colour blue' divides the world into those things which satisfy it, and those which do not, and 'light' functions in more or less the same way. 'Light' is a linguistic device for dividing a range of things in the world, though, in (C'), it divides those things which first satisfy the predicate expression 'colour blue'.

In saying this about the adjective 'light', we are, on the one hand, showing the dependence of that word on another predicate, and, on the other hand, showing it to function nonetheless predicatively.⁵⁷ That these two features can be juxtaposed is crucial. The worry about my examples of alleged predication by objects was precisely their dependency on certain linguistic predicates. But the model provided by 'light' allows us to see that this kind of dependency need be no bar to their fulfilling a predicative (and not merely referential) role. Transposing the account of 'light' to my examples, what I suggest is that we can think of the colour card and swatch as restricting or further subdividing the range of things that fall under relevant linguistic predicates. One can think of the demonstrative 'this' as bringing the colour card into the activity of the sentence, but, once there, the card functions just like an adjective in predicative position. That is to say, the range of things which fall under the linguistic predicate 'colour blue' is, in the context of (C), further divided by the card itself.

Looking at the matter this way allows us to count the colour card as making a fully predicative (or, in my terms, qualificational) contribution, even though there is a dependency on a linguistic predicate. The colour card thus counts as a predicate of the house—it is not merely a referential device—but does so only because the house already falls under the linguistic predicate 'colour blue'.

Accepting the parallel between adjectives like 'light' and the use of swatches and colour samples, it is easy to see why my two examples, though restricted, are nonetheless genuinely cases of objects taking on the role of predicates. Perhaps the

⁵⁶ What I mean is that I am not pausing here to worry about the attributivity of 'light'.

⁵⁷ I said that I would not worry about such things as attributivity, but it is important to see that the dependency at issue owes nothing to this feature of 'light'. Straightforwardly non-attributive adjectives like 'ten-centimetre' are no less dependent on predicates in constructions like 'is a ten-centimetre plate'.

examples are not quite as radical as I first portrayed them—perhaps the colour card and the swatch do not themselves do *all* of the predicative work. Still, even if someone insists that sample-series examples always call on linguistic predicates—perhaps only implicitly—this in no way undermines the point of these examples. For, even if they do not do all of the predicative work, their work is still predicative, or, as I prefer to say, qualificational.⁵⁸

2.8. More and Better Examples

Even accepting that, in sample-series examples, objects do take on roles typically played by predicates, by themselves these examples are unlikely to convince anyone of the pervasiveness of qualification. For that we need examples that do not either invoke the special conventions typical of sample series, or show what could be thought a suspicious dependence on linguistic prediction. In effect, what is needed to counteract scepticism about qualification are cases in which the predicative use of non-word objects matches in generality the potential of objects to serve as referring devices. For the fact is that pretty well any object can be used, in the right general context, to refer to any other, and is able to do so without being parasitic on the referential capacity of names (or similar) in natural language.

What follows is a range of examples intended to convince a sceptic that qualification really does match reference: in the right circumstances, any object, where this includes events, states of affairs, and the like, can be pressed into service to fulfil a predicative function. Moreover, though each example is highly specific, each points to a whole range of further, and different, cases of the phenomenon. (As I am unwilling to discount the swatch and colour card examples, the numbering continues from these two cases.)

3. Imagine being shown a scene of a deserted beach, fringed by palm trees, where golden sands meet a turquoise sea, under a cloudless sky. Imagine further that in the immediate foreground of the scene—at its very focus—is a rubbish bin containing dozens of wristwatches. The rubbish bin and its contents convey a message—give us information about, for example, the simplicity of life in a place like that shown. Indeed, that was the point of this particular advertisement for holidays on a certain Caribbean island. In my terminology, the bin acts as a qualifier; it is an object which we take as providing information about what life is like on that island. Of course, some bit of text

⁵⁸ The analogy with 'light' does however have consequences for our understanding of Goodman's notion of exemplification. In particular, it suggests that he was wrong to think of exemplification as merely a species of reference. In fact, exemplification is only purely referential in those cases where we appeal to an object solely for the purposes of 'concretizing' some predicate. Thus, when I ask my assistant for, say, an augur bit, and am given nothing but a blank look as a response, I might dig such a bit out of my tool box, hold it up, and thereby use it to exemplify (refer to) the predicate 'augur bit'. But Goodman has something richer than this in mind. He surely has in mind cases in which the object doesn't merely concretize the predicate, but rather enhances it. And, as I have argued, we can only understand these cases not as involving merely back reference to the predicate, but as involving predication itself.

could have provided similar information. That is, there are certainly predicate expressions, such as ‘is carefree’ or ‘makes few demands’, which put into words something of the message of the scene. But this in no way detracts from the point of the example. Think here of a parallel argument that might be used in the case of reference. Though we could use the expression ‘my recently bought car’ instead of a salt cellar in telling the sad story of my recent accident, this does not undermine the simple truth that, on the dinner table, the salt cellar refers to the car. Similarly, the fact that we could have used various linguistic predicates instead of the bin of watches does not stop us thinking of the bin as bringing information to bear, and thus as functioning predicatively.

It should be clear that if we had used linguistic expressions instead of, or alongside, the bin, it would simply be wrong to regard the bin as exemplifying any of them. In so far as they are adequate, these expressions will have to be true of the way of life on the island, but they are scarcely true of the bin. And, looking at the converse that Goodman sees as crucial to exemplification, it is simply implausible to count the bin as referring back to any of these predicates. When we ‘get’ the advertisement, we take the bin in the scene, not as referring, and therefore not as exemplifying, but as itself wordlessly doing what could have been done in words. Where linguistic expressions would count as predicates of a lifestyle, the bin counts as a qualifier of just that same thing.

Though each may be obvious, there are two further points to be made about this kind of case. The first concerns the fact that the object has been carefully selected by some advertising agent to get across some message. What I have said so far about qualification might make it seem as if I believe objects to have some kind of mystical power which, in the right circumstances, allows them to speak to us. Since some of the examples yet to come lend themselves even more to this idea of object mysticism, it is important to disown it right at the start.⁵⁹ There is obviously something about the bin of wristwatches that makes it apt for the use to which it is put, but a whole network of intentions surrounding this use is crucial for its success. In this, the use of objects as words is not very different from the use of words as words: communication using one or the other relies on certain intrinsic features as well as on the background and indeed foreground intentions of their users. In the case of words, the intrinsic features—the features words bring to their context of use—consist of what I lazily called their meanings; whereas in the case of objects the story about intrinsic features is more complicated or, perhaps more accurately, more diffuse. For example, in the present case, a proper account of the informational potential of the bin of watches would allude to quite complex social and conventional

⁵⁹ I am hesitant about disowning it completely. The idea of objects speaking is a whimsical way of putting it, but it does call to mind something important: the relation between speakers and audiences. The idea of objects ‘speaking’ only makes sense when there are beings capable of ‘listening’ to them—beings who can, as it were, view the world semantically. (Something like this also applies to referring uses of objects: beings without the relevant semantic capacity could never grasp that special connection between objects we think of as reference.)

features of this item, for example, that rubbish bins are for things rejected, not merely discarded; that watches tell the time but also reflect a certain kind of obsession with it, etc. I shall have occasion in the next chapter to say considerably more about this sort of thing, but here it is enough to note that appeal to some such ‘cultural’ background is necessary (in varying degrees) to the qualificational capacity of objects.

The second point concerns the fact that this example conjures up objects by means of an image; what is suggested is an advertising hoarding or poster of the beach rather than the beach itself. However, this apparent detour through images, or in later examples through my descriptions of various scenes, should be treated merely as an artefact of exposition. In each case, I ask you to imagine that you are confronted, not with an image or my description, but rather with the objects themselves.

4. Qualification—though obviously not thought of in this terminology—is ubiquitous in advertising. But the phenomenon itself has a far wider range. To illustrate this with a rather different sort of case, consider a journalist's comment about the newly rebuilt Hayden Planetarium building in New York: The transparency of the building makes a clear statement about the accessibility of scientific inquiry and our faith in the future.⁶⁰

To get the proper perspective on the example, imagine that, never having heard the above comment, you are walking away from the building and looking back when the sight of the building strikes you forcibly: let us say that you do in fact happen to think some such thought as that above. The large transparent cube, containing the spherical auditorium, conveys to you the thought that scientific inquiry, at least of the astronomical variety, is accessible and thoroughly rational. In my terms, the building itself qualifies scientific inquiry; and, in doing so, it functions very much like a linguistic predicate, though of course the building manages this wordlessly. It should again be obvious that what is involved here is neither just a temporary substitute for a predicate expression we might apply to the building, nor an exemplifying reference to that expression. The thought you have is that scientific inquiry is thoroughly rational, not that the building exemplifies this rationality. (The building might of course be thoroughly rational, but this would be in a wholly different sense from that relevant to the envisaged use as a qualifier. For example, we might think of a building as thoroughly rational if the design optimizes some architectural constraint, such as use of available internal space or its external surroundings, and the planetarium building might do just these things. In this circumstance, the building would be thoroughly rational, and it might even be counted as exemplifying this idea, but this notion of rationality would not be the one required for making the original point about science.)

⁶⁰ Quote from the architect of the Hayden Planetarium, James Stewart Polshek, in a piece about the building by Jonathan Glancey in the *Guardian* (8 May 2000, p. 13 of the G2 section). The accompanying picture shows a huge glass cube containing the spherical auditorium.

Exemplification is simply not of relevance to this example, nor to any others in this section. So, there is no possibility of thinking that my notion of qualification is Goodman's exemplification in disguise. Moreover, the examples in this section cannot be thought dependent on language in the way that was at least intelligible in respect of the sample-series examples. Though those examples didn't explicitly call on linguistic predication, it is possible to imagine such predication at work behind the scenes. Thus, in the colour card example, I could imagine someone insisting that, though no linguistic predicate actually made an appearance, what was conveyed was: Bedroom 2 is *this* colour, where the demonstrative—in indicating the colour card—leans on the predicate expression 'colour'. As already discussed, even if we agree to this, there is no reason to abandon the idea that qualification is at work. For, as I argued in the previous section, the way to understand 'leaning' here is in terms of predication. However, when it comes to the examples in this section, there is no need to sidestep the worry about dependence on linguistic predication, because there is no way to recast them so as to parallel the colour card example. Try it with the example under discussion. Let us say that you have the idea described above about scientific inquiry. With the planetarium building as the target of a demonstrative, you think: It (scientific inquiry) is *that* ... Is there a predicate expression that could be put in the place reserved by the dots, as was the case with the house and '*this* ... [colour]'? I cannot imagine one. (One could have thought: 'it is that rational'. But, as the above discussion of rationality and buildings suggests, this would not be the same thing at all.) The plain fact is that neither the planetarium example, nor any of the others in this section, can be portrayed as leaning on, much less as parasitic on, linguistic predication.

5. The next example reveals a further range of commonplace, though often unnoticed, cases of qualification. Jones is out walking in the country, and trying to figure out what to do about the fact that members of a committee of which she is chair do not agree with her proposed solution to a particular problem. It seems so obvious to her that she is right, and that there cannot be any other way of proceeding. Indeed, so certain is she of her opinion that she is now considering resigning if she doesn't get her way, even though this might harm her standing in the company. The wind had been blowing strongly but is coming on to gale force, and Jones's desire to get home has now overtaken obsessive thoughts about the committee. As she heads back, she notices a small tree on the right of the path that is swaying madly in the gale and to the left a much larger tree standing still and tall. 'That's me!' she thinks to herself looking at the tall tree. Then some way further on, her path is blocked by a yet another tall tree which has only just fallen in the gale-force wind. Seeing this, she decides that she has been foolish: nothing good can come of inflexibly standing

up to the weight of contrary opinion in the committee, and she feels slightly embarrassed by her earlier opinion.

Jones's case typifies many others: objects in nature (where, as I have already said, this includes events and states of affairs) seem often to 'tell' us things about ourselves. Horticulturalists, in particular, tend to find their gardens to be an endless source of wisdom about all aspects of life. My suggestion is that, when we are in this way informed by natural objects and events, these are further cases of qualification. Useful here might be a reminder of something I said earlier: do not confuse the process of qualification with claims such as 'clouds mean rain' or 'red sky at night means that a good day will follow'. These are instances of what Grice called 'natural meaning', but qualification is not natural meaning. Broadly, cases of natural meaning are those in which causal regularities, by their very regularity, provide useful information. Grice of course distinguished these from genuine cases of communication—cases of 'non-natural meaning'—but the natural/non-natural distinction has little to do with the present point. While it is clear enough that tall tree felled by the gale communicated something to Jones, this was neither a case of information gleaned from known regularities, nor, even more improbably, a case in which the gale non-naturally meant something. It is simply that Jones took the event as predicative, as qualifying her behaviour in respect of the committee.

This last observation helps to locate the Jones case on a kind of continuum with the two previous examples. In the advertising case, the creator of the image intended the qualificational upshot described; the object at the focal point of the scene was quite specifically designed to achieve that effect in a viewer, even though the achievement also depends on features of the relevant object. In the planetarium case, the viewer's recognizing the building as informative seems not to depend in the same way on anyone's intentions. Still, given that the building was designed by an architect who intended some such informational linkage, one might think that intentions do play some role, albeit less directly. Both of these cases, then, contrast with the present case: no one arranged for the tree to be felled by the gale, so that the information that Jones takes from this circumstance depends only on her being able to view the world, as I put it earlier, semantically.

The contrast with the advertising case looks sharp but there really is a continuum here. Jones used the event involving the fallen tree to tell herself something—she had the communicative intentions of a diarist, though it was objects and not words that she employed. In typical advertising cases, someone other than the target 'hearer' chose the object and did so with the intention of communicating. However, we can always imagine the qualifying object in fact chosen by the advertiser as just happening to be there. Indeed, in the most subtle instances of visual advertising, communicative intentions tend to be disguised, and we the viewers are invited to do, or think we are doing, the work, rather as Jones did. (The Caribbean advert was not subtle, but even here one could imagine qualification without an overlay of others' intentions; perhaps the beach on the island has a place one can safely leave a watch while swimming. Seeing the repository with these watches, and imagining them as discarded, you take this to contain information about life on the island.)

6. A final literary example comes from Nabokov's autobiography. (There are dozens of examples in it, but I was particularly struck by this one.) In attempting to capture something exquisitely particular about his youthful struggles with the Russian verse form, the *young* Nabokov convinced himself that his lines had a transparency that would convey even the finest details of flower and tree to the reader. However, the *older* writer of *Speak Memory* sees the lines as displaying all too obviously the incoherent effects of the verse form that came, as he put it, 'not by a free act of one's will but by the faded ribbon of tradition'. Helping us to understand what he means, Nabokov describes something he once saw: Years later in the squalid suburb of a foreign town, I remember seeing a paling, the boards of which had been brought from some other place where they had been used, apparently, as the enclosure of an itinerant circus. Animals had been painted on it by a versatile barker; but whoever had removed the boards, and then knocked them together again, must have been blind or insane, for now the fence showed only disjointed parts of animals (some of them moreover, upside down)—a tawny haunch, a zebra's head, the leg of an elephant. (Nabokov 1966: 172) Nabokov says not a word more by way of connecting the comically melancholy fence with his earnest first attempt at poetry; in fact, the above description brings his commentary on versifying to an end. But nothing more needs to be *said*: the more we imagine what he saw, the more we can understand how, for Nabokov, the paling qualified the early poetry.⁶¹

The upshot of these examples, and the many related ones they evoke, is this: we do in fact use objects in the role normally reserved for predicative expressions in natural language. Moreover, the examples suggest that the phenomenon is not something recondite, but rather is familiar and ubiquitous. That it is easy to overlook—and I think it has been pretty well lost to view—is because we think of predication within what is fundamentally a referential framework; it is this framework that, in subtle ways, imposes something essentially linguistic on predication. I have tried to free us from this in part by coining the term 'qualification' but, more importantly, by inviting you to look closely at examples. Though my coinage might be faulted, these examples show the phenomenon to be real enough. Nor, when you think about it, is it surprising that there should be some such phenomenon: reference is not a word-specific function, so there is no good reason for the function of its partner, predication, to be so either.

2.9. Qualification and Predication (Again)

The sample-series examples in section 2.7, restricted though they are, are nonetheless cases in which objects play a role usually taken by linguistic predicates. The

⁶¹ One might insist that Nabokov uses the paling to convey something *to us*, and that this case is therefore too close to metaphor to offer independent support for my notion of qualification, especially given that qualification will figure in my later account of metaphor. But in fact Nabokov tells the story about the paling, not to communicate something about his poetry to the reader, but to tell us how *he* came to understand that poetry. The right perspective then involves seeing Nabokov as having had an experience something like Jones's, but instead of me having to describe that experience to serve as an example, he conveniently did this himself.

examples in section 2.8 show objects as taking on the role of predicates, but also as playing this role without the restrictions appropriate to the examples in section 2.7. What the examples of 2.8 show is that, given the right circumstances and context, any spatio-temporal particular, state of affairs, event, process—items I group together under the label ‘objects’—can bring information to bear on other objects, and thereby fulfil the function that linguistic predication, in its way, also fulfils.

Still, it is possible to imagine a residual scepticism about qualification. In particular, I could imagine someone being suspicious of the fact that circumstances and context play such a large part in it. One thought might be: cases of qualification may well be ubiquitous, but the fact that they require so much stage-setting suggests that they are not somehow central to what we take predication to be. Or another thought might run: unrestricted cases of qualification may not be parasitic on specific linguistic predicates, but qualification itself is only intelligible because of our grasp of the properly linguistic notion of predication.

These two grounds for scepticism are subtly different, but they take as their starting point the obvious fact that we can only understand what is going on in a case of qualification when it is set in a relatively rich context. Nor does the fact that context also plays a crucial part in linguistic predication allay this worry. For someone might well feel that it is the very richness of the context needed for qualification that is the problem. The thought might be: if you tell a complicated enough story about an example, someone might admit that what is going on is very much like what happens in predication. But this is not yet to admit that what are, in the story, called ‘qualifying objects’ are themselves predicates.

Dealing with this form of scepticism about qualification is not easy. On the one hand, it is tempting to insist that it is just wrong-headed. The effort expended in showing that reference and predication are not merely complementary in role, but equal in status, was partly intended to head off this kind of worry. I could with some justice point out that we do not regard the familiar referential use of objects as parasitic on the notion of linguistic reference, and therefore as in some sense second class. Though in the normal run of things, cases of object reference are set in richer contexts than cases of linguistic reference, we surely have reference itself in both cases. So why shouldn't we allow the same latitude to predication, albeit under the less misleading label ‘qualification’?

As it happens, I think this is a perfectly good answer to the sceptic about qualification, but the notion is too important to the central aim of this book to leave it at that. I want even the most sceptical to be prepared to give house-room to my notion of qualification. So, since simple confrontation is not a real option, I turn now to ‘the other hand’.

On this other hand, it might be expected that I would offer yet more examples of qualification—examples which, in leaning less strongly on context, would leave no doubt at all that qualification was not merely ubiquitous but central, or at least fully on a par with linguistic predication. Unfortunately, I cannot do this. There is a kind of case—a range of examples—which would show qualification to be more widespread than even the examples so far given suggest. But the kind of examples I have

in mind, unlike those given so far, are imaginary; they only appear in a somewhat speculative or reconstructive story about linguistic predication itself. Hence, appealing to them is not straightforward. Still, since the reason they are imaginary is directly relevant to the issue of qualification, I intend the whole of the story I am about to tell—the reconstruction, as well as the reason that reconstruction is necessary—to count as a more definitive answer to the sceptic. (There are further reasons for indulging in the storytelling which follows: it will allow me, in section 2.10, to tie up some loose ends in respect of predicates, properties, and Strawson's claim about concepts as 'creatures of language'.)

The reconstruction I am about to embark on begins as a 'just-so' story about how linguistic predication could have come about. So, in telling it, I have to begin by allowing the possibility of a community of human beings who lack any natural language, and I can imagine that this will not play well amongst those who think of language and thought as inseparable. Counting myself in this group—albeit with reservations about what is meant by 'language' in this context—there are several things that I should say to keep those with this conviction about language and thought on side.

First, and most important, my just-so story is intended to be even less ambitious than such stories usually are. When I am finished, I hope you will see certain of our ideas of categorization and predication differently, but I do not expect you to take seriously the background details of the story that lead to these ideas. Second, while the community in the story is said not to employ devices of natural language, they do have the abilities to use objects both as referential and as qualificational tools. That is, though they do not categorize them in the rich sense we associate with linguistic predication, they can track them sufficiently well for the purposes of reference, and also appeal to them as sources of information in primitive kinds of qualification. In fact, one such primitive case will be the centrepiece of the story. Third, while I do see that the idea of a community of human beings who lack language strains credulity, it might help to think of the community I am describing as based on Swift's Academy of Lagado. In Swift's story the members of the Academy gave up using natural language because of their fear that it would mislead them. So, perhaps one could imagine the community I shall describe as a later offshoot of Swift's: its inhabitants continue the languageless tradition of the Academy, but the reasons for this have been lost in the mists of time. That said, there is something importantly different about Swift's Academy of Lagado and my community. His Academicians gave up using natural language, and communicated instead with sacks of objects that they carried on their backs. However, a careful reading of his story shows that these objects are used in what is an exclusively referential way; there is no recognition that objects may also be used predicationally. This is unsurprising, since Swift was essentially satirizing philosophers like Locke, and he thus simply imported the philosophical bias in favour of reference that I have already discussed. Still, as we know, reference alone is not enough to constitute a language: in addition, we require some means of expressing thoughts, and without *both* the ingredients of the basic combination—reference and predication—this requirement cannot

be met. My earlier insistence on the fundamental interdependence of the trio of reference, predication, and truth must be kept firmly in view here. This is why my community has to be seen as able to use objects both referentially and predicationally: without the one, one simply doesn't have the other, and both are essential for truth. (You will be unsurprised to find that I have also dropped the requirement that these objects are carried in sacks.)

Against this background, let me describe a very simple kind of case: Aman has a ewe, and he has never come across another.⁶² (I leave the issue of zoological plausibility on one side, and to make it easy for us, I'll call the ewe 'Clio'.) One day Beman turns up leading another ewe (which I shall call 'Dido'). Aman is surprised: never thinking about it much, he regarded Clio simply as a category-less item in his environment. (Imagine here how you might interact with something in the dark basement corner of an old farmhouse: you walk past and around it from time to time, so you can certainly keep track of it in some sense, but you don't think of it even as a piece of junk, because you simply don't think of it.) However, when Aman sees Dido, her very presence conveys some information to him about Clio; he comes to have a kind of insight as a result of the information that he sees as being brought to bear on Clio. My suggestion is that we can say that Aman comes to see Dido as qualifying Clio. What is that information? Here we have to be careful, as there are difficulties in any of the most obvious ways of spelling it out. On the one hand, it is simply too quick to say that Aman comes to see that Clio is a ewe. This is something we are working towards, but we shouldn't begin by crediting Aman with possession of this concept. For if Aman possessed the concept ewe, it might be felt that he already possessed all that would be needed for employing the predicate 'ewe'. On the other hand, it is no better to say that Aman notices the *similarity* between Dido and Clio. There are indefinitely many respects in which the two ewes are similar—most irrelevant to the issue at hand—and, though there is one respect which is highly pertinent, it would be unduly hasty to credit Aman with having grasped it. For us the salient respect in which Dido and Clio are similar is that both are ewes—they both fall under this (or a related) concept or kind. But, as before, we do not want to begin by treating Aman's insight as straightaway calling on the concept *ewe*.⁶³

What I am insisting is that we must take very seriously Aman's naïve viewpoint: he doesn't regard Clio as a ewe, or indeed as anything other than an item in his environment with which he has interacted in various ways. For this reason, we should not be so ready to say, when he sees Dido, that he comes to understand that Clio is a ewe. Still, it doesn't seem satisfactory simply to keep repeating that, for Aman, Dido qualifies Clio, because, given possible scepticism about qualification, there

⁶² I do not say 'simplest' here and though you will see that the case of qualification about to be described is much more basic than those in the earlier examples, I shall hint at a still simpler case at the very end of the book.

⁶³ I could say much more about the uselessness of appeals to similarity here, but Goodman 1972 has done a lot of the work already. Attempts to call on similarity when in a tight corner in regard to categorization are perennial. Indeed, as will be discussed in later chapters, they figure prominently in certain accounts of metaphor, and I shall therefore have to return to this topic.

does seem a real need to say something more. What I suggest is that we allow ourselves to say, in the above case, that:(I) Aman finds that Clio is Dido-ish.

These are of course our words—Aman has as yet no use for them—and, while I do not mean it to be a replacement for my talk of qualification, in this simple case it can give the flavour of the information that comes with such qualification. Using (I) to characterize the information that becomes available to Aman, we avoid any suggestion that Aman has, or has suddenly acquired, the concept *ewe*. Also, while it is easy enough for us to understand (I), there is something satisfyingly unspecific about it. Indeed, it is just what one would expect of someone who is as conceptually naïve (at least about ewes) as Aman has been portrayed, but who has gained insight from his encounter with Dido.⁶⁴ Finally, (I) should not be confused with the thought *that Clio is similar to Dido*. Similarity is a reflexive relation, qualification isn't: Aman learns something about Clio, not something about Dido, or about both creatures. (Which is not to deny they are similar, or even to deny that Aman might come to appreciate the similarity. It is merely to insist that (I) is not itself a similarity judgement.)

As an aside, let me expand on a previous example which might make clearer the kind of dawning insight that the encounter with Dido produced in Aman. The case takes a little stage-setting, but, since the function of qualification in the Aman–Clio case is so important, I hope you will agree that it is worth it.

Suppose that, having got quite used to seeing and avoiding, though not categorizing, the rusting pile of metal in the basement of your recently purchased old farmhouse, you visit a local farm museum. Encountering there a single metal structure prominently displayed in a glass case, you suddenly see differently—can now construe—what you before only regarded as a pile of stuff in your basement. Even without reading the museum label describing the use to which structures like these were once put, you feel that what was before simply a pile—an obstacle to be avoided—and which you had not thought of as any kind of thing, has suddenly become a something. What I want to say about this case is that the museum object qualifies the original; that it tells you something about it. What it conveys—the information it brings to bear on the object in your basement—is, at least initially, minimal. Not having read the label on the glass case, you don't know what the museum object is, or what it was used for, so you can scarcely be said to have recognized your own pile of metal as a thing of that specific kind. However, it is perfectly reasonable for you to describe your insight by saying that you now recognize the basement object to be a *that-ish* thing (where the 'that' refers to the museum object). As noted, this information is minimal, but it is information, and it strikes you as such.

⁶⁴ It is crucial to the whole of the story as it unfolds that you do your best to imagine things from Aman's point of view. He did not single Clio out as a creature of a certain kind, until he had the flash of insight that came with his seeing Dido. Though I do see that it is difficult to imagine such a point of view amongst adult, social human beings, I am asking you to overlook this difficulty for the further purposes of what I began by admitting is a 'just-so' story.

The perennial temptation with a case like this is to reach for similarity, insisting that what is recognized is that the object in the basement is similar to the one in the museum. However, while this similarity judgement is certainly true, it is simply no help in describing, much less explaining, the insight typical of this kind of case. Aside from the fact that the two objects are similar in all sorts of ways that are just irrelevant to that insight, the plain truth is that, before you visited the museum, you didn't even think of the pile of metal in your basement *as a single object*—an object that could be similar to or different from another such object. It was only after, and as a result of, your encounter with the museum object that you were in a position even to entertain a similarity judgement. An analogy might help to highlight the important point at issue here.

Invited to look through a friend's telescope, you see a blurry white patch surrounding a dark centre. Perfectly reasonably, you say that you don't see anything. Your friend then turns the focus knob one way or the other and you suddenly do see, say, the crescent Venus. What I suggest is that the object in the museum acts like the focusing of the telescope: where once there was an indeterminate pile of metal, seeing the object in the museum brings that pile into focus. Where initially you didn't see (read: conceive of) the pile as anything, now you do.⁶⁵

Let's now move the story on. Suppose it happens that ewes proliferate; that over time there are any number of ewes and owners. Moreover, just as Dido had qualified Clio, we now find that pretty much any single ewe can be counted as qualifying any other. One could say that in this community ewes form a mutually qualifying population. This is of course one way to describe a condition that is necessary for the development of the concept of *ewe*, as well as the predicate 'ewe', but we must try to remain as faithful as we can to the naïve languageless viewpoint that obtains in the population of ewe-owners. All that we can say about this population is that its members notice that any given ewe could be said to qualify any other.⁶⁶ So, while from our point of view, there is certainly a commonality to the insights in each instance of qualification, a commonality that we would expect to be captured by the predicate 'ewe', members of this community lack the resources to express this insight, or even to have a focus for it.

Or at least they lack it until someone comes up with this idea: why not choose a single ewe to serve as the qualifier of all the others? In a way, this would be like an Academician in Lagado using one of his sack of objects—let us say that it is an A—as the standard by which to judge whether some candidate item was A-ish.⁶⁷ Assuming that this good idea is implemented, and that a single animal is chosen in some appropriate ceremony, what in effect has happened is that the community

⁶⁵ It would be good if the farmhouse example could completely replace the just-so story in making out my claim about qualification. However, as I believe that qualification underlies the development of categorization and hence predication, it is necessary, even at the risk of paradox, to describe a more primitive situation than figures in the farmhouse case.

⁶⁶ The temptation to say that they find the creatures similar can be overwhelming, but I continue to demand that it should be resisted for now.

⁶⁷ It also is reminiscent of Platonic accounts of predication, but that is a long story that need not be told here.

now has a Standard Ewe, somewhat in the way that we have standards of various measures. However, one must be cautious about too direct a comparison with such things as standard metre rods. Our standards of length, colour, measure, etc. are parasitic on linguistic predication: we know that something is a certain *length*, *colour*, or *volume*, and the standard is used as a way of keeping our estimations of these already understood notions in order. In contrast, the Standard Ewe is not a way of deciding, for example, how good a ewe some candidate animal is; it does not play the role a standard might for a judge at a ewe-show. Instead, the Standard Ewe is a single selected object which informs the community about other objects, and in so doing is responsible for the beginnings of categorization, not for its refinement. (One might profitably consider the Standard Ewe in the context of prototype accounts of categorization, but the similarities and differences would hold up discussion here unnecessarily.)

The idea of a Standard Ewe has a lot to be said for it, but it comes with certain inconveniences. It is not always easy to bring the Standard Ewe to the place where it is needed, and of course it must itself be replaced from time to time. Fortunately, some bright spark comes up with a further innovation that eliminates these inconveniences. The suggestion is to fashion some object that is relatively portable and durable, and which can serve as a substitute for the Standard Ewe. Taking the idea further, someone in the community produces a little wooden carving of a ewe, and this then serves in the practice in the community in place of the live animal that had been the Standard.

Replacing the Standard Ewe with the wooden model is simple enough, but care is needed to understand exactly what is involved in the transition. As I tell the story, the Standard Ewe is the item chosen to play a role in qualification: when information is sought about some beast, the Standard Ewe is brought out to see if it does or does not qualify it. (One could imagine there being a whole barnyard full of Standards for different kinds, but perhaps the story is elaborate enough already.) In effect, what is sought is precisely the kind of then unanticipated insight that Aman had when he encountered Dido, only with the passage of time, such insight comes to be expected, and is provided by a single Standard. Given this background, it is easy to jump to the conclusion that the wooden model plays this same qualificational role in the categorizing practices of the community. But this would be mistaken. The model itself does no qualifying work, its job is simply that of referring—directing attention—to the Standard Ewe. To be sure, this referential relationship is put in place by there being some kind of depiction of the Standard Ewe by the model, but it is still the Standard that is seen as the source of the categorizing insights. To emphasize this difference between the model's function and that of the Standard, one can imagine that the wooden model itself becomes damaged over time, losing thereby its supposed resemblance to the Standard, but referring to it nonetheless. In such a case, we would scarcely consider the model, having lost its capacity to depict a ewe, as itself the qualifier. Admittedly, you might wonder how the community manages, since, though there is a pointer to it, the categorizing Standard is no longer around. But this is not a serious problem. Memories of the Standard, jogged of

course by the fact that the ewes form a mutually qualifying population, suffices to keep the practice on the rails.⁶⁸

Now while less troublesome than the live Standard itself, the wooden model also has its inconveniences. And it is at this point that a major innovation is made. Since the point of the wooden model was simply to refer to the Standard, someone comes to see that this same reference could be secured more conveniently if, instead of using a concrete object, the community could call on some dependably recognizable vocalization. After all, such a vocalization would always be available, in a way that the model is not. Moreover, just to add a small note of further plausibility to this development, it is decided (or merely just develops) that the vocalization that comes to replace the wooden model is a recognizable ingredient in the noises that were typically made when, in the previous stage, someone went off to get the Standard.

With vocalizations playing the role just described, we are on the threshold of word-use. To keep things simple (from our point of view), I shall imagine the relevant vocalization to be 'ewe'. Strictly, what we might as well now call this 'word' refers to the Standard Ewe, and, as I have been insisting, it is the latter which should be seen as providing information about candidate animals. Of course, by this stage there is no Standard Ewe around, but it still remains true that qualificational information comes not from the newly coined word, but from a conception of the Standard Ewe that remains even after the actual Standard has disappeared.

That said, as the practice develops, it is easy to imagine that this point about the role of the Standard Ewe is gradually forgotten, and that members of the community come to think of 'ewe' as referring, not to the Standard in its qualificational use, but to the information that this use once conveyed. (I shall return to this distinction.) Moreover, this is made more likely if 'ewe' becomes embedded in a now-developing linguistic apparatus that includes the copula. Thus, where in the beginning, members of the community wheeled out the Standard Ewe in order to inform themselves about the nature of some candidate animal, they now say of the candidate that it 'is a ewe'. The phrase 'is a ewe' comes to be a substitute for the original qualificational insight, perhaps even an improvement on it, since the 'is a' construction makes explicit something that was only implicit in the use of the Standard. Thus, whereas one had to understand from the context that the Standard Ewe qualified another *individual*, and not a pair, trio, etc., of such individuals, this construction makes the number of place-holders or 'slots' explicit. (Of course, strictly, there are no slots in 'ewe'—in my story it refers to the Standard, and the issue of slots only arises when the Standard is used qualificationally. But there is a tendency on our part to think of 'ewe' as itself a predicate, and thus as itself having a single slot. That this tendency is natural is clear enough, but it is still a mistake—a mistake that will be explored below.)

⁶⁸ For the purposes of my story, nothing much more needs to be said about keeping the practice going without the Standard, but even this little bit might sound provocative to someone who has taken seriously Wittgenstein's rule-following considerations. However, I think that, so far from contravening those considerations, given space, my story could actually be made to illustrate them.

It might be thought that the introduction of the 'is a' construction at this point violates the ground rules in the background of my just-so story. After all, wasn't the idea to begin with a community that had no use for words? Two things can be said to defuse this worry: first, while exposition is easier if I use our standard copulative construction, we don't have to imagine that something as sophisticated as this really does appear all of a sudden; a simpler indicator of copulation would do, one that would fit well with the primitiveness of developing categories. And, second, I never really meant the story to be one tracing out the origin of language *tout court*. That repeatedly discredited, albeit enjoyable, enterprise is well beyond my remit. Recall that I encouraged the thought of my community as an embellishment of Swift's tale of Lagado. The original Academicians in that story did have language, but they suppressed the use of it, and we could think of their successors as rediscovering the importance of words rather than as developing the whole of language *ab initio*.

Anyway, as already noted, you should not be obsessed with these sorts of detail. The main point of the story is to show, on the one hand, that qualification can be seen as playing a vital part in the development of full-blooded linguistic predication; and, on the other hand, that its figuring only at the beginning of the story explains why qualification tends not to be noticed. These observations are crucial, so let me expand on them.

The typical philosopher's story about categorization takes as its starting point our noticing similarities amongst members of a certain population, say ewes. Without saying that this story is just wrong, I do think that it is not the right place to begin. Aman's insight about Clio did not consist in his noticing that Clio and Dido are similar, but rather that the one qualified the other. Nor are these two descriptions mere variations on one another: aside from anything else, similarity is a symmetric relation whereas qualification is not. To say that A is qualified by B is tantamount to saying something of the form 'a is F', and, in the normal run of things, this does not imply the nonsensical 'F is a'. Moreover, aside from the reality of this difference, I think that judgements about qualification, and not those of similarity, afford us a deeper view of categorization.

To see what I am getting at, return to one of my earlier examples. Walking in the forest, Jones came to see a tree's having fallen in high wind as giving her information about her situation at work. One might be tempted to say that she noticed the similarity between the two situations, but this is not accurate. It is highly unlikely that Jones would have thought of her situation at work as giving her any information about the fallen tree. Yet, as suggested, similarity is symmetric, so similarity cannot be what Jones notices. Someone might try to evade this consequence by insisting that Jones's situation at work is indeed similar to the tree's having fallen, but that Jones only notices the similarity in one direction. The claim might be that *similarity* is a symmetric relation, but *noticing a similarity* is more complicated and can accommodate asymmetries. But this doesn't work. Similarity just is symmetric, so when someone notices—noticing being a factive attitude if ever there was one—that A is similar to B, they must also notice

that B is similar to A. So, whatever it is that is noticed by Jones cannot be the similarity of the two situations.⁶⁹

While the above argument is surely right, I do nonetheless understand the temptation to say that Jones noticed a similarity. This is because we say this sort of thing when describing how it comes about that things are assigned to this or that category—we say that our concept of X is the concept of things similar to each other in appropriate respects. Since similarity is taken as the basal notion in categorization, it offers itself as the most likely thing to say, even in a case like Jones's, where categorization isn't really at issue. However, what I suggest is that, even in the case of categorization, we should not reach so readily for the notion of similarity. Aman's reaction to seeing Dido was of a piece with Jones's reaction to the fallen tree: what goes on in both cases is the use of one object to qualify another. And just as it is misleading to say that Jones noticed the similarity between her situation at work and the circumstances in the forest, so it is misleading to say that what Aman notices is the similarity between Clio and Dido. In the case of Jones, the fallen tree figured as a non-word predicate of her situation at work, and for Aman, Dido figured as a non-word predicate of Clio. There is however a major difference between the two cases: Jones would not have regarded her situation at work as a non-word predicate of the tree's having fallen in the high wind; it is bizarre even to entertain this converse of her original thought. Yet there is nothing strange in regarding Dido as a qualifier of Clio: here the converse *qualificational* claim not only makes sense, it is in fact true.

Notice that the difference between the two cases does not arise from the fact that in one, but not the other, we have a similarity judgement. My point is precisely that similarity does not figure *in either case*. Instead, the difference can be located more at the level of the objects themselves: the objects that figure in Jones's case do not qualify each other, whereas in the case of Aman they do. Moreover, thinking beyond Aman's initial encounter to the situation in which there are lots of ewes around, we can say (as I already have) that ewes form a mutually qualifying population. It is this fact *about these creatures* that leads first to the establishment of a Standard Ewe, and from there to the use of models and then words. In short, it is the fact that ewes form a mutually qualifying class that makes categorization in terms of the concept *ewe* possible.

Discussions of categorization usually begin with the requirement that similarities amongst members of some class of things come to be recognized, whereas mine begins with two separate requirements: the recognition of instances of qualification and the recognition that such qualification is, or can be, mutual. In a sense, my account resolves similarity into what I regard as its components. In the case of Jones, and in the other examples offered in section 2.8, there are certainly instances of qualification—certain

⁶⁹ Fogelin 1988 heroically tries to argue against the symmetry of similarity claims. (This is in the context of his dealing with Beardsley's 1962 argument that metaphors cannot be similes because of the symmetry of similarity and the asymmetry of metaphorical comparison.) However, what Fogelin manages to show, though he doesn't realize it, is that we might well *say* that A is like B, but not that B is like A. But 'say' of course creates an intensional context, and even if A is like B *and* B is like A, we might well say the one and not the other. (I shall not here consider the consequences of this for his defence of the comparison view of metaphor. More on the latter in Chs. 4 and 5.)

objects are regarded as qualifying others. But in all of those examples, the second requirement is not met: the objects do not qualify one another. However, in the case of Aman and his community, this second requirement is met, and, precisely because it is, the path is clear for the establishment of the practice involving both the concept *eme* and its attendant predicate. My just-so story was intended only as one way in which this path might be followed, but however it is managed, once the path is traversed, qualification no longer figures. It has long-since become redundant.

With this last important point, we have come back to my defence of qualification. The examples in section 2.8 seemed to me robust enough to show the importance of the notion, but I imagined someone finding them too dependent on background stories, too elaborate, and, in a word, too rich. The trouble with such richness, according to the objector, is that it makes one wonder just how fundamental qualification is. The contrast here is with object reference: this practice is widespread and easily intelligible. I mentioned Swift's Academy, Wittgenstein's Paris courtroom, and the thoroughly mundane use of items on the dinner table to tell a story. These are some instances of reference by objects, but they do not by any means cover the field. Concrete objects are the most natural item we think of in this regard, but we shouldn't forget that, for example, sounds can be used referentially. The use of musical themes in opera and film to indicate that a certain character is about to figure in the action is familiar, and is certainly a case in which a sound is given a referential function. Images too can take on this kind of function. Goodman struggled—I think rightly—to argue that, in depiction, reference should not be seen simply as a consequence of some independent assessment of similarity. Rather than saying that a picture refers to a person because it resembles him—that reference is here parasitic on similarity—Goodman urges us to recognize that depictive similarity is in fact dependent on reference. Still, leaving the contested notion of depiction on one side, there are many other kinds of case in which images take on a referential function. To take a single, but ubiquitous example, think here of corporate logos, where we have reference, but certainly not similarity and depiction.

The point about the referential use of objects is precisely that it seems unrestricted both in respect of the kind of object which does the referring, and the kind of object referred to. Against this background, the apparently narrower scope of qualification can thus seem suspicious. Taking off from a thought everyone accepts—that reference and predication are the two pillars of at least the simplest truths—I have insisted that, just as reference operates freely outside language, so should predication. Since there is a something strange about talk of ‘non-linguistic predication’, I coined the term ‘qualification’ to label the predicational work of non-linguistic objects. If, as seems to be the case, qualification does not operate as freely as non-linguistic reference, then isn't this a reason to think that I have, to put it bluntly, exaggerated the importance of this notion?

In spite of what might look like self-indulgent speculation, my story about the development of the concept *eme* provides the materials to answer this question. Accepting that qualification plays the role I have described in getting the categorization story going, one can see why it is less visible than its counterpart notion, namely, non-linguistic reference. For, in a way quite different from its referential

counterpart, qualification has to *compete* with language and, as one might expect, it loses out to this more useful competitor. Thus, while we can *imagine* (perhaps with difficulty) the qualificational insight that Aman had, once Aman and his community come to have the concept *ewe*, and the predicate that goes with it, this particular insight serves no further function. Indeed, it is even lost to view in most versions of the categorization story, given common appeals to some sort of similarity.

In contrast to qualification, the referential use of objects, while it can be seen as in competition with language, has features which allow it, in many cases, to win that competition. The Paris courtroom scenario, the dinner party story, the use of sounds and logos, each display some feature that linguistic reference lacks. Exploiting as they do directly visual or auditory cues, there are many contexts in which reference by objects can more than hold its own with linguistic reference. Though not the whole of it, a large part of the parody in Swift's story arises from the fact that the referential use of objects in the Academy is attended by inconveniences far outweighing the benefits of the words one would need instead to effect reference. If Swift had described a practice where objects, sounds or images were used as referring devices in situations where there was no common language, or in contexts where the use of natural language was difficult or inappropriate, then the Academicians would not have looked so foolish.

For some cases of qualification, similar things can be said about competitive advantage: Jones's seeing the fallen tree made her appreciate something about her situation in a forcefully direct way, one that could not easily, if at all, be replicated by words. In each of the other cases too, there are reasons to prefer the non-linguistic route: swatch and colour samples offer perceptual detail unmatched by words; the planetarium building makes its point succinctly; Nabokov's paling is eloquent (and this word is not out of place here). However, in all of these cases, there is some sort of asymmetry between the qualifying and qualified object, and thus these cases differ radically from that of the Clio, Dido, and the other ewes.⁷⁰ In the case of the ewes, and in all cases where there is a mutually qualifying population, the words which come to compete with qualification are the overwhelming favourites in that competition. Why would we stick with qualification by this or that animal when we have available our concept *ewe* and its associated predicate? Given that our interest in mutually qualifying populations is in effect an interest in the development of categories, and given the importance to us of categorization, large swathes of territory become no-go areas for qualifying practices. And the knock-on

⁷⁰ Ah, you will say, but surely in the case of things like colour cards and swatches, we do have mutually qualifying items. And, in support of this, you might add: the colour on the card and the colour of the room, the fabric of the suit and the fabric of the swatch respectively match each other. The reply is this: the second claim is certainly true—the colours match and so do the fabric types. But the first claim is false: the room doesn't qualify the colour card, nor does the suit qualify the swatch. Remember that the conventions governing these series examples make them closer to linguistic predication than is true of other cases of qualifying objects. The colour card and swatch act as sort of side-kicks of the predicates 'colour' and 'fabric type'. However, if you imagine, for example, that the card was used outside the confines of this conventional practice, then it will be apparent that, while it might well qualify the room, the room—an enclosed space with furniture, windows, etc.—would not be taken as qualifying the card. In a sense, I am admitting that series examples are not ideal, but then I have already said this.

effect of this is to make qualification invisible *even in cases where we should be able to see it*.

With this last point, the final piece of my second line of defence of qualification slots into place. Taking seriously the worry that my examples in section 2.8 do not show qualification to be fundamental, I now think I have grounds for meeting the worries head-on. For, so far from being apologetic about it, my considered view is that qualification should be regarded as *the* general notion—a notion on all fours with reference—and we should count predication as a special case, one in which the qualificational effect is achieved by words or phrases.

So far I have been more timid, using ‘qualification’ merely as a label for the predicative function of objects, and have offered various examples meant to convince you that objects really do function in this way. However, the bolder view seems to me more accurate. To be sure, we don't speak as if predication was merely a species of some more general notion—we don't even have a label for any such general notion. Still, if you accept my story about the development of predication, even in barest outline, then you can understand why we might lack such a label and, as part of understanding this, you also see why the less timid view is the correct one.

Our not having a label for what I have called ‘qualification’ results from the fact that, at the heart of our categorizing practice, where qualification might figure, talk of concepts and the use of words have proven vastly more useful. Since in that practice we can only discern individual acts of qualification by imaginative reconstruction, it is no wonder that the label ‘predication’ has come to figure for us as the general name of the function complementary to reference.

Still, though most of our talk is of predication, and many would consider examples of qualification to be only on the margins of our practices, we should not be misled into thinking that qualification is actually marginal. For when you look closely at predication, and come to see the emptiness of appeals to similarity that are often thought to serve as its basis, it is possible to see predication as itself arising from primitive exercises of qualification. Admittedly, seeing it this way requires imaginative reconstruction, but then again so does any other account of predication. The model notion here, as it has been all along, is reference. Without claiming that our referring use of words really did arise from the replacement by words of our use of objects as referring devices, it is at least intelligible that it could have done so. Moreover, if we do imagine that there was some shift from object to word in respect of reference, we find it easy to understand the present relationship between linguistic and non-linguistic reference. My story of Aman and the ewes is intended to fulfil the same purpose in respect of predication. In keeping with the theme of this chapter, that story presupposes the parity between reference and predication as the pillars of the basic propositional combination. Given that the referential function includes both words and objects, so should the predicational function. Having found that the predicational function of objects tends to go unremarked—perhaps not surprisingly, since it is unnamed—I supplied a name and have now suggested a reason for its marginalization. However, if you accept the reason offered, then you can see that its marginalization is, as it were, undeserved. For in giving what I called

‘qualification’ a role in the imaginative reconstruction of predication, I have justified its promotion to the same level of generality as reference.

2.10. Some Consequences

In the course of telling the story of ‘ewe’, I more than once said that we should not look too closely at the details—that it should be taken as merely a just-so story. Indeed, in concluding the previous section, I insisted that what matters most about the story is where it starts and where it ends, not so much the stuff in between. The end point is linguistic predication, for example, phrases of the form ‘is a ewe’, and its starting point is the naïve qualificational recognition that, for example, one ewe could reasonably be thought to qualify another.

Still, while speculations about how one got from the starting point to its culmination in natural language are not offered as historical fact, features of the story as I told it can serve theoretical purposes other than simply supporting the notion of qualification. Moreover, in so far as this is true, it might provide grounds for taking the story a little more seriously, and thus adding further weight to the case for qualification.

First, consider a puzzle in philosophical logic which is not particularly profound, but is vexing. What I have in mind here is the contrast between concepts, conceived of as Frege did, and properties. Many writers talk as if there is no real difference between, for example, the Fregean concept *ewe* (or the ‘universal’ ewe, if this is the way you like to talk) and the property of *ewehood*.⁷¹ Wiggins (1984: 320) finds an obstacle in so thinking. He notes that: ‘the one term we cannot licitly form as a name of [the property of manhood] is: the property of man’. But of course if *man* and *manhood* were interchangeable—if they were both simply ways of specifying some feature or property—this should be a perfectly licit construction. Cutting a not too long story even shorter, what Wiggins claims is that ‘the property of manhood’ is a nominalization of the expression ‘is a man’, and that the concept *man*, specified by the word ‘man’, is a component in this expression. So, it is no wonder that *man* and *manhood* cannot be interchanged intelligibly; they have very different logico-grammatical roles.

Wiggins is certainly right about this, but what he says only labels rather than explains our having notions both of concepts (or universals) and properties. And, short of a kind of collective carelessness on the part of many philosophers, nothing he says explains why it is so common to conflate the two. If, however, you are willing to take seriously the idea that the origin of the word ‘ewe’ (transposing vocabulary back to my story) lies in its standing for, or referring to, something like the Standard Ewe, then it is easy to see why we speak both of *ewe* and *ewehood*, and why there is a tendency to conflate the two.

The Standard Ewe served to qualify any particular ewe: as argued, this is not simply a question of noticing a similarity between two animals, though of course they are

⁷¹ Armstrong takes Frege’s concepts to be his universals and he thinks it unproblematic to take both as properties. See, e.g., Armstrong 1997: 29.

similar, it is rather that the one *informs* us about the other. The Standard Ewe serves the community as a source of insight about this animal, just as, at first, Dido served to give Aman insight about Clio. When the wooden model and eventually the word ‘ewe’, come to be introduced, this practice carries on: the model and then the word ‘ewe’ refer to the Standard, and the latter continues to be the ground of judgements about this or that animal. That said, and as already noted, since there is now no actual Standard around, it would be unsurprising to find a certain confounding of ‘ewe’ as a specifier of the Standard—itsself the source of the relevant information—and ‘ewe’ as a specifier of the information itself. That is, there comes to be little reason in the community to distinguish between the Standard as a qualifier, and the feature or features that are the upshot of such qualification. In effect, there comes to be little to choose between the notions of *ewe* and *ewehood*, since both serve, in suitable linguistic contexts, to get us to the same place. Thus, one can say: Clio is a ewe, or, Clio instantiates *ewehood* (i.e. Clio has the property of being a ewe), without bothering much about the difference between them (except perhaps when taking special care over certain matters in logic and metaphysics).

Still, the two are very different: one can say ‘Clio is a ewe’, but not: ‘Clio is (a) ewehood’. Also, coming at it in the way Wiggins suggested, one can speak of ‘the property of ewehood’, but not of ‘the property of ewe’. However, against the background of my story, one can understand these differences as more than merely grammatical, and as having a clear rationale. To repeat: ‘ewe’ specifies an object which has the capacity to inform us, and ‘ewehood’ specifies, not that object, but instead the information about the world that that object conveys. Moreover, having drawn this distinction, we are in a better position to understand what is right and what is wrong with Strawson's claim that concepts like *ewe* ‘owe their being, such as it is, to words’.⁷²

Strawson contrasts concepts like *ewe* with spatio-temporal particulars. Particulars are said to ‘have their own independent being in the world’, whereas concepts are ‘creatures of language’—a portrayal which leads Strawson to insist that talk of concepts is talk of the senses of expressions. The imaginary story of the Standard Ewe and its transition into full-blown linguistic predication can help us to understand: (i) why Strawson (and others) think this kind of thing; (ii) why there is something not quite right with it; and finally (iii) why there is a deeper (and perhaps even true) *un-Strawsonian* way in which the creatures of language claim can be understood.

2.10.1. Points (i) and (ii)

Dealing with the first two of these points requires no more than that we keep our grip on the separation of function as between reference and predication that I have laboured so hard to keep in focus. The function of referring devices such as names, as Strawson so clearly tells us, is to bring spatio-temporal

⁷² Strawson 1974: 39. These words are in a paragraph that Strawson imagines as put by an objector. But, as noted earlier, it seems fair to attribute them to him, since he says of the objection: ‘let it stand’.

particulars to the attention of an audience.⁷³ Leaving aside Strawson's detailed discussion of the circumstances that are conducive to the coining and use of names, the essential point here is that this function presupposes that there are actually particulars or objects with 'independent being in the world'. How else would what is in effect a kind of pointing be intelligible? However, the function of predication—when this is understood in the general sense that I tried to capture with my notion of qualification—is radically (though complementarily) different. Pointing, bringing spatio-temporal objects to an audience's attention, is not what counts. Admittedly, spatio-temporal particulars do have some role to play. Whether as qualifiers in an unsystematic way (Aman's seeing Dido), or as part of a practice (the story of the Standard Ewe), spatio-temporal objects *serve* predication. Nor is this any less true in the fully linguistic case. As noted at the very beginning of this chapter, words are themselves particulars: in appropriate social and linguistic contexts, the word-object 'ewe' helps us out when we predicate. However, what is crucial is that the function of predication is not like that of reference, and has nothing to do with calling attention to spatio-temporal particulars. That the traditional story speaks of words like 'ewe' specifying the concept *ewe* only serves to confuse the issue. If by 'the concept *ewe*' we mean just 'whatever "ewe" specifies', then of course Strawson is right to insist that concepts are creatures of language. But this is scarcely surprising and is certainly not philosophically significant.

As per the story about qualification that I have told, certain objects happen to inform us about the world. In the most primitive case imagined, Dido informed Aman about Clio. As with names, such cases have a utility that sees them turned into a practice. I imagined the practice beginning with the Standard Ewe—an object which itself qualifies others—but then developing to the point where qualification is no longer intrinsic to the object that figures in the practice. Thus, the carved model succeeds the Standard Ewe, and perhaps it comes to be replaced by some largely non-ewe-ish token. Of course, this exact line of development is not meant to be taken too seriously. Serving more to show a direction, it is the end point that matters—the point at which a word-object displaces imagined tokens and models. For when we see word-objects like 'ewe' as such end-points, we can see how they serve, just as a wooden token, neither as qualifiers nor as predicators, but simply as 'reminders' of whatever does the qualifying (or predicating). For convenience, and in one of its uses, we employ the word 'concept' for what we take ourselves to be reminded of in such cases; we say that 'ewe' specifies the concept *ewe*. Used in this way, words like 'ewe' certainly seem to have referential pretensions, but their referents are certainly not items with the kind of independent existence Strawson attributes to spatio-temporal particulars. But then again it is clear enough that they were never meant to be.

All of this admitted, someone might insist that, in a more Fregean use of the word 'concept', we have a kind of independent existence that Strawson simply overlooked. What I have in mind here are properties. As noted earlier, Strawson does not mention the notion of a property when he makes the 'creatures of language' claim, nor in fact do properties figure anywhere in his book. Yet it might be argued that,

⁷³ The second chapter of Strawson 1974 is devoted to just this task.

on the one hand, ‘ewe’ specifies or refers to the property of being a ewe and, on the other, that such a property has a claim to a real existence independent of language.

Tempting though it is, this line of reasoning only adds to the confusion we get into when we forget the fundamentally different roles that reference and predication play. The word ‘ewe’ simply cannot, within the bounds of grammatical propriety, specify a property, and the imagined developmental story shows why. Using the acceptable expression for referring to a property, we can speak of ‘ewehood’ and wonder whether the property it refers to exists at all and, if so, whether it is independent of language. But just like the earlier use of the concept *ewe*, the property of *ewehood* does not itself undertake the function of predication; we can neither speak of something's being *ewe*, nor of its being *ewehood*. Wiggins rightly suggests that this latter infelicity comes from the fact that something copulative is taken up in the nominalization process that creates ‘ewehood’, but my story suggests a somewhat deeper reason for this. *Ewehood* is not itself predicational, but is rather something like the upshot of a predication's having been exercised. For Aman, Dido qualifies Clio; for the community at large, the Standard Ewe qualifies certain of the animals in the vicinity; and, as we move further away from objects that do the qualifying, the carving and wooden model come to refer to this (or some) genuinely qualifying object. Yet, since in each case of qualification there is some insight gained, there is nothing to prevent the use of yet another object or word to stand, not for any qualifying object, but for the common insight gained when such objects are qualificationally employed. Thus, *ewehood* can be thought of as specifying or referring to what is informationally in common as between this or that qualifying use of an appropriate object. Features or properties, as this common information is more usually called, are thus clearly inappropriate themselves to serve as qualifiers (or predicates), even though they are intimately connected with predicational activity. It may well be the case that properties have independent existence and, in this sense, are more like spatio-temporal particulars, but they are no better than Strawson's concepts when it comes to understanding predication. Perhaps the whole point can be summed up this way: when we seek to discern predication using tools belonging to the world of reference, we are bound to end up either confused or disappointed. We show confusion when we try to think of concepts as genuine objects, or try to think of properties, which certainly seem genuinely independent, as if they were concepts. We show disappointment, as Strawson seems to, when we decide that concepts are merely creatures of language.

2.10.2. Point (iii)

Even granting the reasoning above—reasoning directed to my first two points—there is a more profound, and un-Strawsonian, sense in which we are right to wonder whether even properties are ‘creatures of language’. The following is only the merest hint of how this might work, though it would be a mistake to take brevity here as an admission that what follows is only tangential to the theme of this book. The truth is rather the opposite: what I shall say in the next paragraphs is crucial to my main theme. It is just that the end of this chapter is not the right place to thrash out the issues in detail.

Given my story about qualification, the creatures of language claim is not likely to be true of properties, if by ‘language’ we mean ‘natural language’. Qualification is something that can be done in natural language—there known as ‘predication’—but it is an activity we can engage in without recourse to natural language words and constructions. In so far as properties figure in this activity, it would seem that Strawson's claim that concepts ‘owe their being, such as it is, to words’ simply does not apply to them. What it is about ewes that makes it possible for us to see them as, to coin yet another word, *co-qualifying* is not something put there by our using words like ‘ewe’ and ‘ewehood’. And, though not a standard use, the general description ‘what it is about X that allows us to recognize them as co-qualifying’ can serve perfectly well as a gloss on the notion *property of being an X*.

As so glossed, the notion of a property is independent of words, but it is not independent of our tendency to use objects as qualifiers. We have come to have the words ‘ewe’ and ‘ewehood’ because they codify a range of systematic insights that human beings have had when encountering certain creatures. Of course, everything here depends on unpacking the notion of an ‘insight’. It could (and has been) held by certain philosophers that any insight about categorization worth having is one wholly controlled by properties that are themselves wholly independent of our sensibility. Others are not so sure, allowing, for example, that insights of this sort are either coeval with the very existence of properties, or perhaps are even responsible for them. The underlying point at issue figures in various ways in philosophy of language, philosophy of mind, and metaphysics; it is most certainly not one that can be settled here. However, without even trying to come down firmly on one side, I should like to explain how the very terms of the debate lend themselves to a different and perhaps more defensible interpretation of the ‘creatures of language’ claim.

I think it fair to understand Strawson as maintaining that concepts (and perhaps properties) are creatures of the way in which we construct and understand predicates using familiar words and phrases. In short, they are creatures of *natural* language. Now, while it is true that objects, even when used referentially or qualificationally, are most certainly not words in a natural language, the very fact that they can fulfil these functions gives them some claim to the word ‘language’. A way to put this, one already touched on, is to say that human beings have semantic abilities that are exercised either in the way they manipulate words in natural language or in their taking objects to refer and to qualify. When someone tells a story using objects on the dinner table as props, it is scarcely over-extending our ideas to think this demonstrates a semantic ability. The props function as names, or at least as something close enough to names. If we came across creatures who didn't employ anything we could recognize as natural language, but who could use props in this way, we would certainly think of these creatures as having some sort of linguistic ability. Think about those non-human primates upon whom so many hours of language instruction have been lavished over the years. Even if this instruction were to be successful, no one would think that these chimpanzees have mastered anything as complex as natural language. Yet if it really were the case that chimps could reliably use props referringly (in appropriate circumstances), I think we would

unhesitatingly admit that chimps possessed at least this one fundamental semantic ability.⁷⁴

Of course, cleaving to my earlier remarks about the necessary interdependence of semantic abilities, we shouldn't really expect reference in the complete absence of qualification and truth. An intensively trained chimp might well connect a plastic token with, for example, a banana, and connect it in a way that tempts the thought that the token stands for the banana. But I think we would be right to resist this temptation. Unless the supposed act of reference is accompanied by an ability to use objects as bearers of information about other objects—unless, that is, we have evidence of the ability to entertain (if not communicate) some putative truth—the plastic token shouldn't be taken as referring to the banana; it is merely associated with it, probably causally.⁷⁵

While insisting that linguistic (semantic) capacity, in the broadest sense, shows itself in truth directedness, and that this can be identified only when correlative referential and qualificational functions are fulfilled, let me now focus more particularly on qualification. It was the burden of my story about Aman that the contribution we make to categorization begins with qualification. Aman had an insight about Clio when he saw Dido. It was not that he saw that Clio and Dido shared this or that property or that he noticed that Dido was similar to Clio. It is important that we resist these ways of looking at the matter: concepts and properties are important but they come *later* in the story, and appeals to similarity, while not simply wrong, do not explain what went on. The plain fact is that, before his encounter with Dido, Aman didn't think of Clio as having this or that property or indeed as being separable enough from the background to trigger thoughts of similarity. However, after encountering Dido, Aman came to appreciate Clio in new way: Dido made Clio stand out from the background.

The stages that came after this involved the transition from this initial insight to the full practice of categorization, via the concept *ewe* and its associated predicate. As argued above, these stages also bring properties onto the scene—in this case the property of being a *ewe*, or, more succinctly *ewehood*. Moreover, the qualificational insight which begins the process suggests that the very existence of concepts and properties is dependent on something which is semantic, and is thus a matter of language. This is not Strawson's 'creatures of language' thesis, because the dependence is grounded on something more primitively linguistic than natural language. But it

⁷⁴ The jury of psychologists, once optimistic about the likelihood of non-human primate success in the linguistic arena, shows considerably more pessimism these days. I won't cite references here because the literature is too vast and is easily discovered.

⁷⁵ The interdependence thesis suggests an explanation for pessimism about non-human primate language, even though psychologists do not seem to recognize the fundamentally different contributions that are made to truth by reference and predication nor even, come to that, the fundamental nature of the basic combination. Note too that the parenthetical remark about communication is important. It is no part of the interdependence thesis that a creature adept at qualification uses objects to *communicate* truths to others. Of course, there is a sense in which jointly exercising referential and qualificational abilities so as to entertain a thought about how things are is being ready to communicate thoughts to a receptive audience. But no doubt there are additional social abilities that figure in actual communication, and I think the interdependence thesis shouldn't be compromised by requiring them also to be present.

nonetheless remains true that concepts and properties, in exhibiting this dependence, seem unlike spatio-temporal particulars.

Two important caveats should be entered here. First, I am not myself asserting the dependency above, but merely explaining how, if it were accepted, it shows, in a deeper sense than Strawson's, how it could be that concepts and properties are creatures of language. From my perspective, everything turns on how one is to understand the qualificational insights that I see as grounding the formation of categories and their associated predicates and properties. Such an insight could be an appreciation of something objectively there, waiting to be recognized. On this reading, we discover rather than create concepts and properties. Or we might take the notion of insight to mark the contribution of our shared sensibility, rather than anything there independently of that sensibility. (To the extent that the latter is true, we might have reason to think that concepts and properties are creatures of language in some very strong sense.) Or, finally, it might be the case that there can be a kind of interdependence between 'objective' properties, and the sensibility which is required to appreciate them.⁷⁶

The second caveat concerns Strawson's insistence that spatio-temporal particulars, unlike concepts, are *not* creatures of language. I repeated his claim, but I did so without any conviction. For I lean towards the view that spatio-temporal particulars are no less 'creatures of language', so long as we understand this phrase in the metaphysically non-committing sense just discussed. But this is a long story and one it is not appropriate to tell here.

A final note about Goodman's notion of exemplification to end this section. While I have been at pains to distinguish qualification from exemplification, I can imagine someone still struggling to keep the two notions properly separate. However, the story about how predicates and properties emerge puts the difference between these notions in stark relief. Goodman is always careful to speak of exemplification as of predicates, and not as of properties. This is because of his thoroughgoing nominalism. Now, whether or not you share Goodman's scruples, the very fact that he was concerned about this feature of exemplification shows just how far it is from qualification. For it simply doesn't make any sense to worry whether qualification is as of predicates, or of properties. That a given object qualifies another is something wholly independent of, and in my view prior to, our talk of predicates and properties.

2.11. Summing Up

To those interested in metaphor, but having no particular interest in arcane debates in philosophical logic, my excursion into that subject might well seem overlong. Yet for others, what I have said will seem merely to scratch the surface of a subject that has

⁷⁶ Some idea of the range of views here can be got by comparing the kind of thing David Lewis says about fundamental properties (1994: 412–13) with the broadly Wittgensteinian line about shared practices that emerges from his rule-following considerations. For Wittgenstein, categorization is at least in part dependent on those practices, whereas for Lewis the only kind of categorization worth having is that which reflects a reality of properties—the fundamental ones—wholly independent of anything deserving the label 'our practices'.

been central to contemporary philosophy for more than a hundred years. To both of these constituencies, I give the same answer: everything that I have said in this chapter will in the end prove necessary for the exposition and defence of my view of metaphor.

The central theme of this chapter has been the sharing of functions as between words and objects. If my main interest had been solely in the use of words as names or labels, the chapter would have been less convoluted, and much shorter. Words are used referentially, and so are objects. The way that reference works in natural language is subtle and often deceptive; the use of objects as referential devices is often, though not always, more straightforward. But there can be no doubt that, for example, in explaining an unfortunate accident, a fold mark on a tablecloth can refer to the rue Jacob in Paris just as do the words 'rue Jacob'. However, as luck would have it, my view of metaphor requires the investigation, not so much of reference, but rather of predication. And here I found it impossible to avoid the longer story.

Beginning with certain hints that can be gleaned from Goodman's notion of exemplification, I found it necessary to lift the lid, even if only slightly, on the narrow world of philosophical logic. Scrupulous though it is, I found that, in that world, predication gets a raw deal. Lip-service is paid to its importance—it is held, along with reference, to be essential for our most basic propositional structures, for thoughts themselves. Nonetheless, it is not afforded parity of treatment with reference. Nor, given that predication tends to be explained in a framework dominated by reference, is this inequality even noticed.

My particular complaint was this: reference can be accomplished both by words and, in the right circumstances, by non-word objects, whereas predication is typically thought of as inherently and solely a function of words. This seemed to me to be wrong, and I set out to see whether there was any real ground for this differential treatment. Finding none, I conjured up a label—'qualification'—which functioned exactly at the same level of generality as 'reference'. Moreover, at the risk of being thought whimsical, and the even greater risk of being taken for a semiotician, I set out to show that there is genuine point to this notion. That is, I set out to show that there are cases in which we use non-word objects as, in effect, predicates.

Circumstances have to be right and context must take up some of the slack: amongst other things, we need context to indicate the number of 'argument places' that objects can possess. But there really do seem to be actual examples of objects (including, as always, events, states of affairs, facts, and the like) providing predicational information. Moreover, in a long aside I won't even try to summarize here, I have suggested that there are interesting further reasons why qualification is not noticed—reasons which, once acknowledged, lend further support to this notion.

What I haven't done yet is to use the notion of qualification in my account of metaphor. This is the job for the next chapter. However, I cannot resist ending this one by pointing out that, if I prove right about qualification being fundamental to our understanding of metaphor, this will be yet a further, and substantial, thing to be said in its defence.

3

The Semantic Descent Account

In the previous chapter, I set out to see whether some of things we happen to do with words could also be done without them, but right at the beginning I cautioned against thinking that the results of that investigation could be immediately applied to metaphor. It is certainly true that framing metaphors is something we do with words, but I didn't actually get further than a consideration of the more basic linguistic functions of reference and predication. Nor is metaphor going to be added to *that* list here. For, though the idea of a wordless metaphor is perfectly coherent—indeed it is something that will crop up in Chapter 4—the single item on the agenda of this chapter is an account of metaphor *as a linguistic phenomenon*. Having established that qualification is predication, or is in effect predication, in which objects and not words do the work, we need now to combine that result with the fact that in metaphor things begin with words.

Yet another cautionary note should be inserted here. In what follows, I will set out my account in a fairly minimalist way, saving many elaborations and embellishments for the two chapters which follow. This minimalism will be most evident in the simplicity and paucity of examples. As will I hope become apparent, I have not streamlined my exposition in order illicitly to gain plausibility for the account, though this certainly seems a common enough strategy in the literature.⁷⁷ But in writing this chapter, I have been keen to keep the account itself, and certain aspects of its defence, in sharp focus. The wide range of features typical of metaphor and the variety and richness of examples can easily overwhelm any exposition. So, to repeat, in this chapter I shall keep things simple, defending the account where necessary, but making sure that its overall structure stands out.

That said, as noted in the Introduction, minimalism has its risks: starkly simple examples can make certain kinds of objection seem pressing, even though, against the background of richer examples and further considerations, many of these objections should fall away. Still, as honesty requires it, I shall flag up these objections in this chapter, while nonetheless inviting you to reserve final judgement until you read Chapters 4 and 5.

⁷⁷ Roger White (1996) is especially good at uncovering examples of this, and he is not reticent about denouncing them. In aiming to avoid this himself, he may perhaps err on the side of an over-rich diet of examples, but his castigation of overly simple accounts is certainly a welcome relief from rather one-dimensional examples and selective discussions. There will be a proper discussion of his account in Ch. 5.

3.1. Metaphor and Semantic Descent

The notion of ‘semantic ascent’ is too familiar to need much of an introduction.⁷⁸ Asserting that the sky is blue by saying ‘The sky is blue’ is using language in its most ordinary way. Call this the ‘ground floor’ use of language. Saying of the *sentence* ‘The sky is blue’ that it is true achieves largely the same purpose as the original assertion, but it does so by engaging in a bit of semantic ascent; it is moving up a level of language by speaking, not immediately about the sky itself, but about the truth of words which themselves speak of it. Assertions about the truth of sentences are typically first-floor uses of language, being one level up from the ground (at least in Europe). What I claim goes on in metaphorical utterances is a bit of semantic *descent*, but, perhaps surprisingly, it is descent that begins from the ground floor of language use, and moves down to what we can think of as a sort of ‘basement’. Since I don’t expect any of this to be familiar, I shall discuss a couple of examples, in the course of which it should become clear not only what semantic descent is, but how, in allowing us to tap into qualification, it gives us a fresh way of looking at metaphor.

Consider again Romeo’s assertion: (R) Juliet is the sun. Assuming as I have throughout that (R) has already been identified as metaphorical, the outlines of my proposal are quite simple. When (R) is understood as an ordinary assertion at the ground-floor level of language, it is either false, or perhaps even a bit of nonsense. But I suggest that, instead of pausing over any such reading, we take the metaphoricality of (R) to demand a kind of semantic descent. Instead of thinking of the word ‘sun’ in (R) as a word that plays only its usual natural language role in the predicate ‘is the sun’, think of the object that this word stands in for—think of the sun itself. Both the sun and the word ‘sun’ are objects, albeit of radically different kinds. The one is the fiery nuclear star at the centre of the solar system which supports life on earth; the other a set of marks that play a special role in a complex linguistic practice. Yet focus just on the fact that they are both objects. We do not have to think of the predicate in (R) as simply fulfilling the ground-floor predicate role its word-objects have been assigned in a complex social practice. Think in addition that the function of the word-objects in (R) includes taking us from this ground-floor level to a level below—to the basement-level of non-word-objects. In the specific case of (R), ‘the sun’ is pressed into service within a natural language predicative structure to take us to the sun, and it is the latter object that gives us information about Juliet.⁷⁹ The hearer is invited to understand this object as a qualifier

⁷⁸ I presume that Quine 1970 is the original source for this term.

⁷⁹ I put it this way to remind us that, strictly speaking ‘is the sun’ is the predicate, and ‘the sun’ is only an ingredient in it. As noted in Ch. 2, there is a certain carelessness in usage here. Sometimes philosophers speak of the predicate ‘ewe’ and sometimes of the concept expression ‘ewe’ which figures in the predicate ‘is a ewe’. I think that the latter is correct, but the shorthand is convenient and certainly intelligible. In the present context, it seemed necessary to try to be accurate about this, though as will be shown at length my proposal is not limited to subject-predicate metaphors.

of Juliet in just the way that, in the examples of Chapter 2, swatches, colour cards, dustbins, buildings, fallen trees, and palings qualified their subjects. Hence, though the initial setting is wholly one of natural language, much of the work of (R) is accomplished by means of the not naturally linguistic, though predicative, mechanism of qualification. Since ordinary natural language predication is at the ground-floor level, the move to the level at which qualification figures is a move to a sort of basement level.

Without stopping to discuss this at length here, there is something appealing about the idea that metaphor works in the basement. After all basements are where one finds foundations, and there are those who urge us to see metaphor as somehow implicated in the very foundation of language. Having already suggested that qualification plays such a role in categorization, it is not difficult to see how semantic descent followed by qualification might satisfy that urge.

What I have given so far is only an outline of the semantic descent proposal, and much more needs to be said (an understatement if ever there was one). However, as is perhaps obvious, the sentence used in the Romeo example has special features which can be misleading, so it will be useful to have a second one in play before I consider the proposal in more detail. For that I shall adopt one of Davidson's. Begin by supposing that someone has said: (T) Leo Tolstoy is an infant. Assuming that the background to this assertion makes it unproblematically literal—perhaps it is said on an appropriate date by one of Tolstoy's uncles—then we can understand 'infant' as making its usual linguistic contribution to (T). As Davidson puts it: How is the infant Tolstoy like other infants? The answer comes pat: by exhibiting the property of infancy, that is, leaving out some of the wind, by virtue of being an infant. ... Tolstoy shares with other infants the fact that the predicate 'is an infant' applies to him; given the word 'infant' we have no trouble saying exactly how the infant Tolstoy resembles other infants. ... Such similarity is natural and unsurprising to the extent that familiar ways of grouping objects are tied to the usual meanings of usual words. (Davidson 1984*a*: 247–8)

However, as used by a critic of the adult writer Tolstoy, as in the remark, 'Tolstoy is a great moralizing infant', that Davidson cites, (T) is metaphorical. Having noted this, Davidson mocks the idea that what we should now do is to find out what the set of infants, now including Tolstoy, have in common, thereby stretching the meaning of 'infant' to include this particular adult. The linguistic object 'infant' figures in a complex social practice, serving in that practice to label a property or properties we naturally discern. Nothing is gained by stretching that word-object so that it counts, either generally or on a specific occasion, as a label for a property-complex encompassing both toddlers and the adult writer Tolstoy. But no such stretching is required by the semantic descent story.

Understanding (T), when it is used metaphorically, begins with the fact that the word 'infant' applies to the usual suspects, namely the set of not yet grown up

human beings whose properties encourage us to group them together.⁸⁰ No stretching of meaning, nothing special here. What happens next, however, is that, understanding this perfectly ordinary linguistic label, the hearer moves from it to an exemplar—to one of the set of things to which the predicate ‘infant’ applies—and then this exemplar serves as a qualifier of Tolstoy. It is as if, instead of coming out with (T), the critic said:(1) Tolstoy is this ...while pointing to a 2-year-old child.

Here I expect questions to come crowding in. Since neither (R) nor (T) have explicit demonstratives in them, how are we to understand their relationship to sentences like (1)? Aside from issues surrounding (1), how does one choose an exemplar of a predicate? Does one actually have to have a specific infant in mind, or does any infant count? How can one be sure that the exemplar chosen will serve in the metaphorical setting? That is, how does one control the predicative use of an object? What happens if the exemplar doesn't happen to be, or is not thought to be, of the right sort for the metaphorical predication?

These are of course all reasonable questions, and they will be addressed at some point in this chapter. Basically, they cluster around two issues. On the one hand, there are questions about the movement, the descent, from words to objects—questions about the movement from ‘is the sun’ and ‘is an infant’ to such things as the sun and an infant. On the other, there are questions about the suitability of these objects to fulfil the qualificational role required by my account. The next three sections will concentrate primarily on the first of these issues. This is crucial because, unless I can defend the movement from words to predicative objects in metaphor, all the work put into the notion of qualification, and the work still to come, will be to little purpose. The remainder of the chapter will then concentrate on the issue of qualification itself.

3.2. Metaphorical Predication and Demonstration

First, let's look at:(1) Tolstoy is this ... (*said while pointing to an infant*),and its supposed relationship to the original metaphorical utterance:(T) Tolstoy is an infant.My offering of (1) as a way of thinking about (T) is not intended as a proposal of strict equivalence. Still, it is not stretching things too far to think that the demonstrative

⁸⁰ I am perfectly happy to follow Davidson here in speaking of a kind of sharing of properties as characteristic of categorization. But this is not going back on the discussion of these matters in Ch. 2, because I have been careful to resist saying that categorization is *explained by* or *grounded in* our noticing that properties are shared or that certain things just seem similar.

version captures something about the metaphorical sentence, and that is all I need. For I shall argue that when we understand the demonstrative sentence properly, it contains a clue about the metaphorical (T) which points to the correctness of the semantic descent proposal.

The issue of demonstratives used as in (1) came up in Chapter 2. There I was grappling with the question of whether the swatch and colour card cases—cases I called ‘sample series’—were genuine instances of qualification, or were instead merely cases in which some demonstrated object supplemented or filled out an otherwise linguistically articulated thought. This question arose because some might think that:(2) My house is this blue (*accompanied by pointing to a colour card*),is, as far as the demonstrative is concerned, not much different from:(3) Put the ice in this.This putative similarity is problematic because in (3) we have, as it were, full linguistic articulation: we recognize that complete practical understanding requires one to know the actual referent of ‘this’, something the context surely ought to provide, but there is nothing essentially linguistic missing. If (2) were similar, then this would ruin the point I was trying to make about the predicative-like role I claimed to have found in the demonstrated colour card.

My counter-argument required us to look at sentences such as:(4) My house is light blue.In (4) we have a structure parallel to that in (2), and exploiting this, I claimed that the predicative function of ‘light’ is in fact matched by the predicative function we must now allot to the demonstrated object in (2). Even though the colour card's predicative function is dependent on the linguistic predicate ‘blue’, it nonetheless has such a function. So, we cannot simply dismiss the demonstration in (2) as merely a contextual filling out in the way that seems natural enough in (3). The demonstrated item in (3) is important for a full understanding of the sentence, but doesn't itself have a linguistic function, while the colour card demonstrated in (2) does.

Whatever you think of this argument, I remind you of it again here because I shall not rely on it. Instead, I shall argue that we should distinguish sharply between what is going on in (1), and what is going on in both (2) and (3). When you come to see that the demonstration in (1) involves neither contextual filling out, nor even a subsidiary predicative role, my hope is that you will be prepared to see it as requiring the more radical treatment that comes with my semantic descent proposal.

Superficially, the utterance of (to remind you):(1) Tolstoy is this ... (*while pointing to an infant*),resembles:(2) My house is this blue (*accompanied by pointing to a colour card*).

But, whereas the presence of 'blue' in (2) guides our use of the demonstrated object (the square on the colour card), what is going on in (1) doesn't fit this pattern. This difference is easier to see if we consider a more austere version of (2). Imagine someone uttering: (2') My house is this ... (*while pointing to a square on the colour card*). Though there is here no explicit use of a guiding predicate, one has no trouble at all in finishing this sentence with 'blue' or 'colour'; indeed, (2') seems to call out for some such completion. But try this doing something parallel with (1). Here are two possibilities: (1') Tolstoy is this infant, (1'') Tolstoy is this human being. Neither of these work in the way that the completions of (2') do, nor, in fact, do they advance matters at all. In (2) and in the completed (2'), the demonstrative helps to secure a further narrowing of the general division of things into blue and non-blue, but this is palpably not what is going on in with the demonstrative in (1') or (1''). Moreover, the latter sentences have the same bizarreness as the original metaphorical (T). Saying that Tolstoy is this (demonstrated) infant or this (demonstrated) human being is no improvement on saying (while demonstrating an infant) that Tolstoy is *this*. Nor would it help to insist that what is demonstrated in (1') is the property of being an infant, rather than an infant itself. Partly this is because it is difficult to read (1') as demonstrating a property, but mostly because properties don't help here.

Note first that (1') is only superficially like (2). In the latter, one can be satisfied that what is demonstrated is the property of being a specific colour because, in the end, a property is precisely what (2) attributes to my house. But it makes no more sense—is no help with the metaphor—to say that Tolstoy possesses the property of being an infant than to say that he is an infant. Not only is it wrong to say that Tolstoy is one and the same as *this* (demonstrated infant), as apparently required by surface form of (1), it is no less wrong to say that he has the property of being an infant, or that he has the property of being this particular kind of human being, namely, the infant kind. (I will return below to further consider the role of properties in the demonstrative (1).)

It is important to be clear what is at issue here. As admitted above, I am not claiming that the use of the demonstrative in (1) to refer to a particular infant is simply fine. It isn't. The most straightforward way of reading (1) is bizarre, and some story must be told about what is really going on. But whatever that story is, we must have one. We cannot get away with thinking that (1) is fine because it demonstrates a property or properties, and is thus something like (2) or (2'). Moreover, the very fact that, as it stands, (1) is bizarre means that we cannot treat it as like (3); as like a case in which the demonstration has as its point the supplementation with something extra-linguistic of a perfectly intelligible linguistic construction. The demonstration in (1) is not straightforwardly like that in (3).

Against this background, my suggestion is that (1) is best explained via semantic descent and qualification. As in ordinary sentences with demonstratives like (3), the demonstration in (1) aims squarely at a spatio-temporal particular. Though not without a certain strangeness, (1) requires us to supply an infant as the referent of ‘this’. In so far as (1) captures something of the original metaphorical (T), what we have here is semantic descent. (More on the relation between (T) and (1) below.)⁸¹ As noted, supplying a particular for ‘this’ leads to a certain strangeness: it invites the reading of the copula as an identity claim, as if we are saying, bizarrely, that Tolstoy is one and the same as some infant. However, by calling on qualification here we can overcome this temptation, and, at the same time, remove the strangeness of (1). For unlike the supplied referents in ordinary demonstrative sentences, the object demonstrated in (1) has, in addition to its being a particular in a context, a linguistic function: the infant answering to ‘this’ in (1) qualifies Tolstoy. In effect, the expression ‘is this’ in (1) is functionally a hybrid. It consists of a *word* ‘is’—understood as the predicate copula and not as the sign of identity—and the *object* answering to ‘this’. The copula and the object working together function as a predicate of Tolstoy.

Note the way the copula exerts some control over the qualification effected in (1). The object called upon by ‘this’ wordlessly exercises a predicational function, but this is partly because it is set in a linguistic structure typically marking monadic predication, namely, ‘is (a) ...’. We are thus encouraged to understand the qualifying object as itself ‘monadic’. This helps with a problem that emerged in Chapter 2 in respect of qualification, namely that objects, in contrast to linguistic predicates, lack ‘slots’. In Chapter 2, I insisted that the absence of slots was not itself a reason to be suspicious of the idea of object-predication. What I claimed was that slots indicate, in a fully explicit way, what outside natural language can also be indicated by the circumstances within which the qualification takes place. But, on the proposal which finds a hybrid in (1), we can see how it is possible for the ‘adicity’ of a predication to remain a matter of words, while the predication itself is accomplished by objects and not words. (Of course, I have yet to consider examples in which objects function other than monadically. That will come mostly in the next chapter.)

The fact that the treatment of (1) starts with a fully linguistic construction (‘is this’) and ends up with a hybrid (‘is’ + object) suggests another kind of linguistic control operative in metaphorical sentences, but spelling out this suggestion will take a few paragraphs.

3.2.1. Qualification and linguistic control

The demonstrative ‘this’ invites us to pass with the least possible informational baggage from words to objects; all the work is

⁸¹ More certainly needs to be said about the very idea that we have this kind of reference in (T). Note that the Tolstoy example is one of those mentioned earlier: along with many simple subject-predicate metaphors, it has features of which raise questions about my account which would just not arise in realistically rich and complex cases. Nonetheless, it is useful in other respects, e.g. in this discussion of demonstrative sentences, precisely because of its simplicity.

accomplished by the extra-linguistic circumstances or context within which ‘this’ is used. However, the original focus of my semantic descent proposal: (T) Tolstoy is an infant, did not involve demonstratives. As noted earlier, my interest in the demonstrative in (1) is that it captures something of the import of (T), though I never claimed any *equivalence* between (T) and (1), or between metaphorical utterances of this general form and sentences with demonstratives. For a start, ‘infant’ is not a demonstrative, nor could it be plausibly argued that this concept-expression contains a demonstrative element. Moreover, there are endlessly many metaphorical sentences in which explicit demonstration would be simply out of the question. Remember that ‘objects’, as I am using this notion, includes items such as events, situations, and states of affairs, some of which might well be non-actual in any straightforward empirical sense. (The problem of non-existent objects called on in metaphors will be discussed at the end of this section.) Still, within the context set by (T), it doesn't seem unreasonable to think that (1) at least approximates (T)'s message.

For a start, both sentences are problematic: on the surface, they are bizarrely false or perhaps just plain bizarre. Further, the source of this bizarreness is pretty much the same in both cases: the author of *War and Peace* is not an infant, nor is identical with *this* (demonstrated) infant. Still, if we take the hint that (1) offers, both sentences are made intelligible by my proposal. The sentence with the demonstrative ceases to be bizarre if we take the demonstrated object to be doing qualificational work. And the original sentence (T) likewise comes out alright if we take the concept-expression (‘infant’) as inviting semantic descent to an object that falls under this expression. For, though this descent is not accomplished demonstratively, the making of it nonetheless offers us a way of making sense of (T). Rather than taking it to claim infanthood of Tolstoy, we are free to take it as claiming instead that some infant qualifies him.

In this case of semantic descent, as in all others, we are invited to move from words to objects. Note though that the words from which we descend can play an important guiding role in qualification. Moreover, there is every reason to think that such guidance goes well beyond the fact that the object descended to must fall under the concept delineated by the words in the original sentence. This is the suggestion about linguistic control mentioned earlier.

Obviously enough, the infant used to qualify Tolstoy must be an infant. But there are other concept expressions that have infants in their sights. Thus, we might have been told, for example: (5) Tolstoy is an early stage but independently viable human organism. While it is true that any exemplar of this predicate expression is also an exemplar of ‘infant’, it is of course absurd to think that this sentence is just as good for metaphorical purposes as (T).

It might be thought that this is actually a worry for my semantic descent account of metaphor, but only by someone who had forgotten the lessons of the last chapter,

and is not paying attention to the central aim of this one. Qualification is *not* something that can be guaranteed to work just by wheeling in an object in the presence of some target subject. Context and circumstance are crucial to the intelligibility and aptness, as well as to the usefulness of any instance of qualification. Moreover, and crucially, when it comes to metaphor, we are not dealing with qualification on its own, but with qualification that arises from an encounter with an utterance or inscription in natural language. The words that figure in any such encounter are therefore as much a part of the context of the qualification as is any feature of the object itself. More specifically, it matters a great deal whether the words from which descent is made are as in (T) or as in (5). That is why the inappropriateness of (5), so far from being a problem, is actually a pointer to an important positive feature of my account. The objects reached by semantic descent from the words in each of (T) and (5) may be the same, but the words themselves guide or control or prepare the ways in which we can be understood to use that object to qualify Tolstoy. This is not to say that these words guide us to different objects—that has already been made clear—nor do they encourage us to posit anything as problematic as ‘objects under descriptions’. Instead, in serving themselves as part of the context of utterance, the words exercise some control over the way the object got by descent comes to figure as a qualifier. Thus, infants are infantile, and this latter expression, while it can mean simply ‘pertaining to the early stage of human development’, offers more than a hint of the qualificational role that the object, the infant, is intended to play. In contrast, it is unlikely that the object got by descent from the words, ‘early stage but independently viable human organism’ would be taken to qualify the adult Tolstoy in the same way. This is not because the objects differ, but merely because the explanation of how the objects come to be used depend in part on the words which leads hearers to them.

An aside in two parts: first, the Tolstoy example makes the point about linguistic control seem weaker than it would be in realistically complex examples. As I keep saying, I will return to discuss the downside of this kind of simple subject-predicate example.

The second point is actually a sort of disclaimer. I am aware that, in making the point about linguistic control, and indeed in giving an exposition of the semantic descent account generally, it can sound as though I am propounding a psychological theory; as though what is in question is how we actually process metaphors. However, this should be seen merely as an artefact of the demands of exposition: making a point about how, from a theorist's point of view, one should account for a feature of an utterance tends to make for a great deal more circumlocution than writing as if one was adopting a hearer's point of view. More will be said about this later in this chapter.

3.2.2. Semantic descent and properties

Here let me consider one last issue to round off this section, an issue which requires us to return to the idea that properties are somehow involved in understanding (1) and hence (T). Earlier on, I noted just how

difficult it is to understand (1) as involving demonstrative reference to properties. We have no trouble in understanding the previously discussed: (2) My house is this ... (*pointing to a square on a colour card*), as demonstrating a property of the square, rather than the square itself. If for no other reason, this is shown by the naturalness of adding the property-invoking expression 'colour' to (2). But there is no way to mimic this in the case of (1): adding the expressions 'infant' or 'human being', if it does anything, intensifies the bizarreness of (1) by emphasizing the particularity of the demonstrated object.⁸²

There thus seems to be a real difference between the use of the demonstrative in (1) and in (2); that is, in metaphorical and non-metaphorical contexts. My semantic descent proposal suggested a way in which we could leave untouched the reading of (1)—a reading on which a particular object is demonstrated—while yet rendering the utterance intelligible. This required us to see the process of qualification at work behind (below?) the surface reading of (1). Now someone might see this as an opportunity to insist that, if we were prepared to abandon this surface reading of (1), we could make it intelligible without appeal to anything as radical as qualification. In outline, this would work as follows. First treat (1) as actually saying something like: (6) Tolstoy has some of the properties of this ... (*an infant is demonstrated*) Clearly enough, (6) takes liberties with (1), but if we allow them, doesn't (6) show how to render (1), and ultimately (T), intelligible? That is, isn't it perfectly sensible to explain the metaphor (T) as claiming, not that Tolstoy is in fact an infant, but that certain properties of infants apply to him?

Appeal to properties and the idea of similarity (understood as a sharing of properties) is perennial in discussions of metaphor: even before we got to metaphor, it surfaced in the discussion of qualification, and we will have reason to encounter it again. Here let me point out two interconnected things. First, if we are attentive, we should find (6) no more intelligible than the untampered with (1). Thinking carefully about the properties of infants—remembering all the while what they are properties of—it would be strange to think that Tolstoy has any of them. He does not totter around uncertainly, babble, smile gormlessly when funny faces are made at him, scream when put to bed or complain when denied access to potentially dangerous objects. Nor does he cling to his mother or drool. (I realize I am not painting a flattering picture, but the point here is to take seriously the idea that what is

⁸² Only touched on implicitly earlier, what about completing the demonstrative this way: 'Tolstoy is this kind of thing'? Superficially this sounds alright, but 'blue' and 'kind of thing' are radically different. Subject to vagueness, which is just not relevant here, 'blue' names a property—it effects a division of things into the blue and the non-blue. 'Kind of thing' is simply not like this. On the most plausible reading, saying that Tolstoy is this kind of thing might well just be a general way of saying that he is an infant or a human being. On a somewhat forced reading, it may be taken in the way explicitly given in (6) below. If it is, then I take back my claim that we can only achieve what (6) does by twisting the surface syntax of the original.

in question are properties of infants.) To this the typical riposte would be: no one is claiming that Tolstoy has *just* the properties that an infant has, rather he has properties ... well, *like* those of the infant—he is infant-like. Now, whatever else one thinks about this response, my second point is that this is no advance on our understanding of the metaphor which began the discussion.

Being told that Tolstoy is an infant is simply not the same as being told, for example, that he is a writer. (This is putting it mildly.) Some explanation must be offered of the former remark, an explanation that is not required of the latter. It sounds informative to say of the original (T) that its point is to attribute some of the properties of infants to Tolstoy, but this is an illusion. It may well be that the properties of infants can be thought of in some way—can be transformed—so that we can see them as applying to Tolstoy. I would scarcely deny this. But then the focus of our attention should be on the processes of transformation of properties, and we cannot, as some may be tempted, think of the claim in (6) as a finished job of work.

It will become clearer than it might be at present that some version of the appeal to properties is in direct competition with the semantic descent/qualification account. On my view, the understanding of (T) and (1) requires that we descend from words to an object (an infant), which is then used to qualify Tolstoy. The competition has it that (T) and (1) invoke properties of infants in the attempt to understand what is being said about Tolstoy. My view takes at face value the surface readings of the relevant sentences, but it requires one to accept my story about qualification; the other view requires us to put aside the surface reading, but it works with the familiar idea of a property. So far, it might be thought, too close to call. Not so.

Several things speak in favour of my account and against appeal to properties. However, engaged as I am in trying to make clear what is involved in my account, and especially in the semantic descent part of it, I shall only outline them; more criticism of similarity will follow in Chapters 4 and 5.

First, a property-invoking account should not be thought in genuine competition with mine unless it can cover the same range of examples. Given my present self-denying restriction to simple subject-predicate cases, the competition looks real enough: predicate expressions of the form 'is an F' lend themselves to talk of properties. But when it comes to syntactically more complex cases, it will be difficult even to formulate a property-invoking conception. Nor will the problem be just one of complexity. The richer examples to be considered in later chapters will involve certain phenomena of metaphor—phenomena such as that of deadness in metaphor and mixed metaphor—which are simply not amenable to the property treatment.

Secondly, even with respect to a simple case like (T), the field on which the two views play is far from level. On my account, the move from words to objects is crucial, and, though I haven't emphasized this, the same is actually true of the properties view. For what is in question in (T) is not the property which is characteristic of infants—the property of being one—but rather some property or properties of those things which fall under the concept-expression 'infant'. In other words, there is a kind of implicit semantic descent, or something like it, involved even on the

supposedly competing account: you have to think of particular infants to get at their properties before you can even begin to think of how to transform these properties in ways appropriate to Tolstoy.

Thirdly, the unexplanatoriness of the properties story about metaphor is more fundamental than might appear on the surface. I have yet to give any real detail about the constraints needed to make qualification yield plausible explanations in this or that case of a metaphor. And someone might think that this lack of detail shows that my account and the properties-based one are, at least at this point, tied. But there is a difference: there is reason to think that, in respect of the properties-based account, the route to these details is blocked.

At bottom, the properties account treats (T) as claiming that Tolstoy and an infant share certain properties. As we have seen, it is not that specific features of infants are asserted to be features of Tolstoy. Rather, it is that there are ways of transforming features of infants so that they then can plausibly apply to Tolstoy. How does one go about transforming such a feature? An infant might scream when it doesn't get its way. Tolstoy doesn't. But, if we re-conceive this property of an infant, perhaps we can come up with something that does characterize Tolstoy. For example, suppose that, when certain of his purposes are frustrated, Tolstoy writes some bitter denunciation of the person or circumstance held responsible for his frustration. Can we not say that Tolstoy and the infant share the property of lashing out when thwarted? This kind of property transformation seems natural enough—though it is not always going to be as easy as this—yet there is a problem with it. When we transform properties in this way, so that they become bland enough to apply both to the infant and Tolstoy, we lose sight of what was metaphorical in (T). We have left behind properties specifically of infants, and it was these that made the metaphor apposite in the first place. This suggests a kind of catch-22 for the property account: if you don't leave the actual properties of infants behind, you cannot achieve the transformation required to render (T) intelligible. Yet, if you do leave them behind, you have somehow lost the metaphor. No such problem dogs my account: by insisting that the infant as such qualifies Tolstoy we keep the original metaphor firmly in the picture.

It would be easy to misunderstand my opposition to the property account. It is not that I think it wrong to say, in regard to (T), that Tolstoy shares certain properties with infants. I think this true, and I don't think that any writer on metaphor would deny it. Even Davidson can allow that someone can come to think, as a result of hearing (T), that relevant properties are shared. What I object to is thinking that all we need to do is to advert to some such sharing in order to account for the intelligibility of (T). I have outlined some of the reasons for this, and, as already noted, I will amplify and add to this list in due course.⁸³

⁸³ I have been careful not to mention the views of Josef Stern (1985, 1991, 2000) in my discussion of the demonstrative in (1), though I can imagine that anyone familiar with those views would be puzzled by this omission. Basically, the reason for it is that I want to be fair. His account of metaphor, though it does have more trouble with demonstrative constructions than he seems to realize, seems to me to fall down, not on this relatively technical issue, but on matters more connected with the issue of properties and similarity. I will get to his view in Ch. 5.

Aside from this or that specific objection to the property account of metaphor, there is a more deep-seated reason for my thinking it fundamentally unappealing. The idea of qualification precedes—and at least partially explains—the ways in which we come to speak about concepts or properties. This was the burden of my remarks about the origins of categorization in the previous chapter. Putting these remarks together with my insistence that qualification is a crucial element in understanding metaphor, it should come as no surprise that I regard any appeal to properties as hopelessly too late. They have already been encountered as by-products of the account of qualification—the very notion at the centre of my account of metaphor—so they are unlikely to impress when they are re-encountered in property accounts of metaphor. For me, appealing to properties in order to *explain* metaphor is something like making introductions in a room full of people who know each other already. Your use of their names might be accurate, but whatever you are doing it is not effecting introductions.

3.3. Predicates and Exemplars

In this section, I should like to address some concerns about the examples that I have used. Initially, I shall try to be more explicit about the similarities and differences—especially as concerns semantic descent—between: (R) Juliet is the sun, (T) Tolstoy is an infant. I introduced (T) because I was concerned that someone would think (R) loaded the dice too much in my favour. Romeo's remark, in having both the definite article and a proper name, suggests reference to a particular more directly than sentences of the more 'standard' metaphorical form typified by (T). By appealing to a metaphor with such an obvious referential device, it might have been thought that I was making it too easy for semantic descent. However, (R) is actually stranger than it might appear, and this *is* relevant to my account. Having been careful not to derive any undeserved support from any particular features of (R), I now want to show how these features in fact lend merited support to my account. (Long-held-over issues raised by the descent in cases like (T) will come in the second part of this section.)

3.3.1. Romeo's predication

I begin with some obvious observations. Sentences of the form: *_ is the ..., are by no means semantically uniform. Here are some ways of filling the gaps:* (7) Benjamin Franklin is the inventor of bifocals. (8) Einstein is the brilliant scientist. (9) Ernest is the most awful bore. No doubt I have overlooked many other variations, but the above illustrates the range of possibilities I shall call on. The first asserts an identity by using two

referring expressions, though of course this is subject to issues about the referential status of definite descriptions that are not of concern here. The second has some claim to this same identity status, but it would be naïve to think that that was all there is to it. Unlike (7), it does not simply contain two ways of picking out individuals which are then asserted to be identical. In some way or other, the reference to ‘the brilliant scientist’ conspires with the copula to yield what is in effect a predicate of Einstein. The third sentence is also superficially similar to (7), but, like (8), doesn't seem to be simply an identity claim; it is impossible to avoid the feeling that the point of the copula here too is fundamentally predicational. Yet the superlative in the (9) adds an interesting extra element.

Suppose that, instead of (9), someone had asserted: (10) Ernest is an awful bore. Ernest might not think so, but this is a weaker condemnation than (9). Moreover, this is pretty much what one would expect: changing the definite to the indefinite article seems bound to have some weakening effect. However, consider the claim: (11) Ernest is a most awful bore. Though containing the indefinite article, those I have asked find this sentence to have pretty much the same strength as the original (9). There is, though, something strange about it: while my informants found (11) perfectly idiomatic, the superlative ‘most’ and the indefinite article do not really go together. This is perhaps clearer in a case where the superlative form of an adjective does not rely on outside help from ‘most’. For example, consider the distinctly odd: (12) Ernest is a fastest runner. What this suggests is that there is some pressure to write ‘the’ in sentences like (9), even though they are intelligible with the indefinite article, and make the same point as their definite article versions. It is as if our desire to point out the extremity of Ernest's boringness leads us to the hyperbole of ‘most’, and this in turn puts pressure on us—grammatically though not necessarily semantically—to use a definite article.

The various points made about sentences (7) to (9) will come in handy in dealing with Romeo's (R). However, before I come to that, note a respect in which his sentence differs from any of them. The definite article in: (R) Juliet is the sun, is semantically (though certainly not grammatically) redundant: the expression ‘the sun’ is one of those cases where the definite article, and what follows it, form a semantically unitary, though typographically complex, proper name. There might be some resistance to this observation because of the tendency of non-astronomers to use ‘sun’ to mean ‘star similar to the sun in having orbiting planets’, or even sometimes just ‘star’. But, leaving aside the fact that Shakespeare's sentence would not have been understood this way, and the fact that it is strictly incorrect astronomical

usage, we can avoid arid controversy by the simple expedient of capitalizing. (In fact, some writers do this anyway.) Thus, from now on think of Romeo as having said that Juliet is the Sun; surely, in this sentence the definite article is redundant in the way described.

Now, given that (R) involves reference to a unique, actually existing particular, it would seem that (7) is the sentence we should look to; we should see the copula in (R) as that of identity. However, this is most certainly not how this sentence is taken, and this is somewhat mysterious. Many writers have taken Romeo's sentence as a good example of a metaphor, not least because it has a kind of vividness absent in the usual 'Harry is a fox, wolf, tiger ...' sorts of case so common in philosophical treatments of metaphor. But when this sentence is held up as a reasonable example of a metaphor, the fact that it looks like an identity claim is not even noticed. Though I haven't done a head count, I cannot think of any writer who treats (R) as other than a better (more vivid) case of a subject-predicate metaphor. In an effort to shed some light on this, let me return now to sentence (8).

As noted above, the expression 'the brilliant scientist' is not understood in (8) as simply referring to a specific individual who turns out to be Einstein. Only someone impervious to the nuances of language would hear (8) as just like (7). Instead, we understand (8) as *somehow* using the superficially referential expression, 'the brilliant scientist', in a more purely predicative role.⁸⁴ How might we cash the vague 'somehow'? Well, suppose the expression 'the brilliant scientist' directs our thinking towards an exemplar, towards an individual, not necessarily a historical figure, who is quintessentially a brilliant scientist. And suppose further that, in focusing on this exemplar, we come to think about Einstein in a particular, or even new, light. (I am imagining that (8) is used in a context where it is informative and not merely emphatic. Think of someone saying it to students in a high-school physics class who are perhaps too young to have much of an idea of who Einstein was and is.) What we have here is of course a form of semantic descent followed by a process like that of qualification, and though there might be other stories that one could tell about what is going on in (8), I can't think of any that so naturally explain how what looks like straightforward reference to an individual could be turned into something predicational.⁸⁵ (Note though, that I stop short here of saying that the exemplar here works in precisely the way that it does in metaphors. More on this important issue nearer the end of this section.)

The minor mystery about Romeo's remark that was described earlier disappears completely if one sees it as like (8): the descriptive name—'the Sun'—does indeed refer, but the sentence is not naturally understood as like the claim that

⁸⁴ It might be thought that (8) should be heard as: 'Einstein is *the* brilliant scientist', where stress plays a semantic role. However, I think that the stressed version is yet another variation, and its availability does not take away from the fact that we have predication and not identity in (8).

⁸⁵ Perhaps it is not necessary to say this, but just in case: many identity claims are informative. This is something we knew even before Frege made us realize its importance. But what we have in (8) is not an informative identity, but rather a sentence that seems to be an identity, but is in fact predicational through and through.

Benjamin Franklin is the inventor of bifocals. Instead, we take the referent as itself having a *predicational* (or better, given that an object does the work, *qualificational*) effect, and that is why (R) is assimilated to the usual range of subject-predicate examples of metaphor. Someone could I suppose insist that (R) does assert an identity whose patent absurdity alerts us to its metaphorical character. But this is just not the way it tends to be taken. Romeo is not informing us that the world contains an object which happens both to be Juliet and the Sun—he is not identifying her with the Sun—though he is certainly using the Sun to tell us something about her. Any patent absurdity in his remark belongs with the ‘usual’ metaphorical absurdity we find in claims about a man Harry that he is, for example, a fox, or a wolf or a lion. (Note, though, that I am not here signing myself up to the idea that patent absurdity or falsehood is necessary to the identification of metaphor.)

Suppose instead of (R) that Romeo had said:(R') Juliet is a Sun.Is (R') even intelligible? I ask this because, given that there is only one thing that answers to ‘Sun’, the existential ‘a’ might be thought inappropriate. However, a little informal canvassing has convinced me that (R') is not only intelligible, it is in fact acceptable, indeed it is almost as strong as Romeo's original remark. In fact, the relationship between (R) and (R') is very much like that between (9) and its existential variant. So that you can take them all in, they are:(R) Juliet is the Sun.(R') Juliet is a Sun.(9) Ernest is the most awful bore.(11) Ernest is a most awful bore.The slight difference in strength between the first and second of each pair (and the hint of oddity in the second sentence in each pair) has precisely the same cause. ‘Most’ in the second pair implies uniqueness, as does the capitalized ‘Sun’, so it can seem odd to team them up with the uncommittedly general ‘a’. Yet when we find ‘the’ instead of ‘a’, this doesn't itself inform us of the uniqueness of ‘most’ and ‘Sun’—we knew of this uniqueness already, so in both cases ‘the’ is strictly redundant—but it is still reasonable to think of the uniqueness as *reinforced*. This is perhaps the reason why the second of each pair is heard by some as slightly weaker than the first, though others find them merely stylistic variants.

I have spent a lot of time on the nuances of the definite and indefinite articles (in a single natural language) because I believe them to show something important about subject-predicate metaphors. More particularly, what they show seems to me to constitute further support for the semantic descent account. Here is how the story goes.

The sentence:(13) Harry is an accountant,

seems to be syntactically and semantically of the same form as: (14) Harry is a fox. But there is more going on in (14) than is usually recognized. In whatever way we manage it, let us suppose that (14), on some specific occasion of its use, is identified as metaphorical. (As I keep saying, my account of metaphor does not take any line of the issue of identification, but we can surely agree that, even if it is not particularly vivid, (14) is a metaphor.) Having identified (14) as metaphorical, we certainly realize that, whatever else is going on, we will not be able to explain adequately the linguistic act effected in its utterance by treating it as just like (13). However, as a prelude to any heavy-duty theorizing (including my own), consider: (14') Harry is the fox. This sentence might be less conventional, but it is surely just as good as (14) for conveying the metaphor. Now, as per previous discussion, there are two ways to read the copula in (14'). On the one hand, we can see it as inviting an identity between Harry and some specific fox, rather as in 'Benjamin Franklin is the inventor of bifocals'. On the other, we can see it as predicational, rather as in 'Einstein is the brilliant scientist' or 'Ernest is the most awful bore'. Given the implausibility of the first reading, the second is certainly indicated, and quite predictably it clears the way for my preferred account. That is, given that the expression 'the fox' is most naturally taken as referring to an exemplar—a specific yet typical instance—of the predicate expression 'fox', we have to find some way to combine this reference to an exemplar with the predicational aspirations of the 'is' in (14'). And this is precisely a situation that semantic descent and qualification were designed to accommodate. In referring to an exemplar, we semantically descend from the words 'the fox' to a relevant object, while preserving the predicational aspirations of 'is' by treating this object as itself a predicate—as a qualifier of Harry. In the way described earlier, we are forced to recognize that (14') contains a hybrid predicate, one composed of the word 'is' and a non-word object. (Remember too that we cannot get away with seeing 'the fox' as contributing the straightforward property of *being a fox* to the sentence, since this is certainly not something (14') attributes to Harry.)

What about the humdrum: (13) Harry is an accountant? Everything about the argument of the previous paragraph depends on the fact that meaning, and metaphoricality, are preserved in the shift from the indefinite article version in (14) to the definite article version in (14'). If the same shift can be as easily and conservatively effected in the case of the standard subject-predicate (13), this would create problems. For surely it is implausible to see semantic descent and qualification at work everywhere. As it happens, though, this is not a problem I will have to face, as the shift in (13) leads to: (13') Harry is the accountant,

and the most natural interpretation of this calls on the identity reading of the copula, not the predicational one; it is thus quite different from (13).⁸⁶ Nor can this difference be attributed to the fact that ‘fox’ is a predicate expression for a natural kind, whereas ‘accountant’ is not. (I mention this because someone might think that reference to an exemplar is typical of natural kind terms, and only of them.) For a start, ‘infant’ is not such a term, though the article shift also works here. Moreover, a little reflection shows that the article shift works just as well when the terms in play are least like natural kinds. Try it with the metaphorical: ‘Ruth is a (the) bulldozer’; and the literal: ‘The yellow vehicle parked in the field is a (the) bulldozer.’)

Though simple subject-predicate metaphors are a highly restricted category, the fact is that they differ in the way just described from simple but literal subject-predicate sentences, and this difference is congenial to the semantic descent account. In effect, ‘standard’ metaphorical subject-predicate sentences like (T) or (14) have more in common with the non-standard (R) than might at first have seemed plausible.

In the next chapter, I will consider metaphors of widely different syntactical forms, showing in fact how, in dealing with these complications, my account gains even more support. Still, it is worth having spent some time on the subject-predicate form because, aside from its prevalence in philosophical discussions, it can seem *least* suited to my account. This is because my account requires us to find some object to which we are referentially directed by some plausible understanding of the words used in the metaphor, but the *indefinite article 1 noun* structure typical of subject-predicate cases does not immediately suggest reference to some object. I earlier attempted to get around this by calling on a demonstrative version of the Tolstoy example, and this might have done the trick for some readers. But parity between: (14) Harry is a fox, and: (14') Harry is the fox, is intended to offer further support. I claim that, however much subject-predicate metaphors resemble ordinary subject-predicate sentences—sentences innocent of object-reference—they are not. They more closely resemble the sentences: (8) Einstein is the brilliant scientist, (9) Ernest is the most awful bore,

⁸⁶ I suppose the stressed, ‘Harry is *the* accountant’ might just achieve predicational status, but it also introduces further complications not relevant to the present point. Note too that it is not necessarily even a consequence of (139) that Harry is an accountant. Though I don't want to get involved here in tangential issues about attributive and referential descriptions, the utterance of ‘Harry is the accountant, but he is not an accountant’, sounds fine to me, and it doesn't take much imagination to see why.

than they do non-metaphorical sentences of the form ‘A is a B’. And, while (8) and (9) are not themselves metaphorical, these sentences contain descriptive phrases which are at least candidates for reference.

Nonetheless, they are not metaphors, so does the semantic descent account apply to them and, if so, does this invite further worries about the promiscuity of that account? The answers to these questions support rather than threaten my account. On the one hand, it does seem as if something like semantic descent is a part of our understanding of these sentences. Each invites us to think of an object—the brilliant scientist/the most awful bore—and thus there is the suggestion of a move from words to objects that is characteristic of semantic descent.

On the other hand, the *role* of the object got by this semantic descent is somewhat different from the relevant counterpart in metaphorical contexts. Thus, having thought of an object answering to the description, ‘the brilliant scientist’, and regarding this object as a qualifier of Einstein, we can nonetheless spell this out in a way not available in the metaphorical case. Thus, we can treat features of the envisaged brilliant scientist as attributable *directly* to Einstein: this is part of what makes the sentence ‘Einstein is the brilliant scientist’ a *literal* predication. No such possibility exists in the case of Juliet and the sun, or Tolstoy and some particular infant. As already noted, there is a temptation to think that the property-sharing story can explain these metaphors, but, as I have argued, it doesn’t, though, as the Einstein example suggests, property-sharing might well be all we need in non-metaphorical cases of descent and qualification. If the brilliant scientist is someone who is careless about practical matters, but able to make huge imaginative leaps in the attempt to explain the natural world, then we understand (8) as claiming that these features apply to Einstein. Indeed, any property of the brilliant scientist is intelligible when applied to Einstein, though of course there can be arguments about which ones are most salient or true. In contrast, think of features of the sun and Juliet: the sun is a nuclear furnace, it is responsible for the earth’s warmth and light, it is used (was used) to measure time, etc. These can scarcely be attributed to Juliet without calling on some story of property transformation which is not easy to tell, and which would in any case undercut the explanatory pretensions of the appeal to property-sharing. Moreover, as I have earlier argued, in telling a story about how properties of the sun might be ‘sanitized’ sufficiently to be attributable to Juliet, we would have lost sight of the original formulation that we were trying to explain.

An interim conclusion: the sentence Romeo used is generally taken as a subject-predicate metaphor, but someone might think that, since Romeo’s sentence contains a referential expression in its predicate, my use of this example unfairly softens one up for my semantic descent account. In fact, it is much more complicated: for a start Romeo’s remark is not strictly of the subject-predicate form. However, when we look at this sentence more closely, and compare it with sentences like (8) and (9) above, it is possible to expect a reader to be genuinely (and fairly) moved in my direction. Admittedly, I am only basing this conclusion on the evidence of some pretty hair-splitting comments on a range of English-language sentences. But this evidence is

better than nothing, and I mean you to add this to what I have given so far and also to what is yet to come.⁸⁷

3.3.2. Tolstoy and other infants

Most of the effort of section 3.3.1 was expended on the Romeo example, but there are more things to be said about the one involving Tolstoy. As I have several times acknowledged, there is bound to be a certain resistance to the idea that (T) involves semantic descent to a determinate infant, or even to an exemplar. Does ‘infant’ really involve reference to a determinate individual? Do we really have a particular infant in mind? There are several strands to my response to these questions.

On the one hand, it is important to remember that my account of metaphor is not intended as an account of how speakers and hearers process metaphors. Though it is often convenient in my exposition to talk about what goes on when someone encounters a metaphor, semantic descent is not in the end meant as a description of any such actual ratiocination. I am thus not claiming that we hear (T) and mentally work out which, if any, determinate infant is in question. Instead, the semantic descent account attempts to characterize metaphors in such a way as to make their intelligibility possible: the thought is: *were* we to allow descent from ‘infant’ to a determinate individual, and *were* we to imagine this individual taking on a role usually played by a predicate—a role I called ‘qualification’—then we could make sense of (T).

This distinction in perspectives should be familiar from the literature about theories of meaning for non-metaphorical utterances. On the model I think we should favour, characterizing what is generally thought of as the meaning of a word or phrase is part of the project of making sense of utterances involving that word or phrase. It is not a psychological account of what speakers and hearers actually know, though it is an account which suggests that, if they knew the relevant meanings, we would find their interchanges intelligible. (Actual speakers and hearers might process utterances in ways that do not directly require them to know the meanings described by the theory; the theory simply shows how the intelligibility of utterances comes to be possible.⁸⁸)

⁸⁷ For those who do not find hair-splitting problematic, I add a speculative remark about the reason many do not even notice the ‘ungrammaticality’ of: ‘Ernest is a most awful bore’ or ‘Juliet is a Sun’. There is some ground for thinking that the natural language marker of predication in the simplest subject-predicate sentences is the expression ‘is a’ and not simply the ‘is’. If this is right, then it is not surprising that we find these sentences perfectly alright: it is as if one were saying: ‘Ernest is-a (the) most awful bore’ and ‘Juliet is-a (the) Sun’. Construed in this way, ‘(the) most awful bore’ and ‘(the) Sun’—these very things—become predicates in virtue of their being marked as such by ‘is-a’, and it is therefore not surprising that we do not find real tension between the indefinite article and their definiteness. It scarcely needs saying just how congenial this is to my account.

⁸⁸ In this paragraph I have spoken of meaning and meanings, but this is just convenient shorthand, and is most certainly not meant to exclude the Davidsonian truth-conditional account. It is especially important to note this because—whatever some of his supporters have written—Davidson is firm in adopting the non-psychologist stance described in the text.

Long experience of trying to explain this distinction has convinced me just how difficult it is to grasp, never mind defend, but if you are willing to adopt the perspective on offer, it should lessen somewhat the oddity of treating ‘infant’ in (T) as calling on a determinate infant or exemplar. It is not that I am insisting that a hearer struggles to bring this individual to mind as part of comprehending (T). Hence, it is no objection to my account that you cannot find yourself engaging in any such mental activity when you hear this sentence. The claim is rather that, in adopting a theorist's point of view, assigning an individual to ‘infant’ is a way to make (T) intelligible.

All that said, honesty requires me to admit that, aside from its being controversial, the distinction of perspectives is not by itself enough to allay worries about the descent from ‘infant’ to an infant. So, a second defensive strand is called for, and it is at this point that I need to be explicit about the ways in which simple examples like (T) can mislead us into having unnecessary worries about semantic descent.

What is obvious about (T), and about virtually all similarly simple subject-predicate metaphors, is that they tend to be, at the least, tired, or, even more often, dead. A useful hint that they are like this comes from the fact that many dictionaries actually list a meaning for ‘infant’, which would in fact extend to the adult Tolstoy, and all dictionaries include such a meaning for the closely related ‘infantile’. Now the very fact that simple subject-predicate metaphors are dead, or close to it, does not by itself help me out in respect of descent. If (T) is a dead metaphor, and many think that this means it is in some sense no longer a metaphor, why did I use it? And even if it is only tired, how does this affect the issue of semantic descent?

As these questions show, appealing to the energetic status of (T) to allay worries about semantic descent would seem to require deeper investigation of the general consequences of that status, and that is something not attempted until Chapter 4. But even without taking on the whole of this task here, there is a way forward.

Begin by assuming or even just suspecting simply this much: precisely because (T) and its simple subject-predicate cousins tend to be tired or even dead, less is required from any account of metaphor for their intelligibility. In respect of my account, this means that there is less pressure to use the full resources of semantic descent to make (T) intelligible; what is known about ‘infant’, something shown in any reasonable dictionary, removes some of the need to descend to a determinate object, which then serves as a qualifier, in order to make the metaphor intelligible. This is of course only an assumption, but it gains support, even in advance of a full appreciation of the phenomenon of dead metaphor, by imagining ways to inject some energy into tired or dead metaphors like (T). Instead of the simplest subject-predicate form that we have in (T), think of: Tolstoy is a infant who has just had his favourite toy taken away, or: Tolstoy is an infant who cannot get the attention of his parents.

I am not saying that these are metaphors worthy of great literature, but they are certainly more vivid than the original. Crucially, moreover, they are also cases in which we are less inclined to think of what is in the dictionary, and more open to entertaining the idea that a particular infant figures in our assessment of Tolstoy. And this is precisely what one would expect, given the way the semantic descent account is meant to work.

The above consideration is most certainly not intended as direct support for the semantic descent account; it is intended only to defuse an objection which is, I think, an artefact of the simplest kind of subject-predicate example. Useful though they are for allowing me to sketch my account with a minimum of fuss, their very simplicity can be problematic. A full appreciation of these problems will have to wait on the discussions in Chapter 4 both of complexity and of the phenomenon of dead metaphor. Still, the fact that one can breathe life into a tired subject-predicate metaphor precisely by making the predicate more likely to fix on some individual should make more plausible the idea that a determinate object plays a role in the full account of its intelligibility.

Given the above considerations focusing attention on the way objects come into an understanding of metaphor, I can no longer ignore a problem that has been ticking away in the background. I begin with some comments about actuality.

My account might seem to depend on there always being some *actual* object that is in the extension of, or that serves as the exemplar of, the relevant predicate expression in a metaphor. But this is simply not true. The fact that all the cases of qualification in Chapter 2 involved actually existing items is perhaps partly responsible for this misleading impression, as is the unhelpful fact that the Romeo example uses the sun, an actual object if any is. However, it shouldn't be surprising that all pre-linguistic, pre-metaphorical cases of qualification involve actual objects, since the only kind of confrontation one could have with an object in these cases is perceptual. However, in spite of the Romeo example, this is not typically the case. Metaphor is something we do with words, and it has been my contention that we can best understand what is done when we take some of the words in a metaphor as confronting us with an object. But this kind of confrontation is quite different from that involved in perception. Words like 'infant' or, in a more plausibly vivid case, 'infant who has just had his favourite toy taken away' apply to things, have extensions, and the first part of the task of making metaphors intelligible requires us to conceive of some determinate infant which lies in some such extension. Conceiving, perhaps even imagining, some such infant is what is required by semantic descent, but there is no requirement here that the object be actual, that it be something we could see, touch, or pass from hand to hand.⁸⁹

⁸⁹ Semantic descent is clearly a species of reference, albeit one quite different from other members of this large genus. I will say some more about this and about how we can conceive of a determinate, even if non-actual, object in Ch. 4. Also, in that chapter I will finally put to use the conception of an object I have insisted upon throughout—a conception on which events, facts, states of affairs, etc. also count as objects. I mention this here because there should be even less temptation to be worried about the issue of actuality when 'object' is interpreted this widely.

It is precisely at this point that the worry alluded to above can no longer remain in the background, and it can be summed up in a simple question: what about unicorns? Philosophers who read my comments about conceiving of a determinate member of the extension of some predicate will instantly wonder how there could be semantic descent when a predicate is, as is said, empty. There are infants, and it might be allowed that metaphor requires us to conceive of being confronted by one of them in the sense required, but what about a metaphor built around, for example, words like ‘unicorn’? My answer to this question can be fairly brief.

The kind of empty predicates that are in question here belong to fiction in a broad sense, and fiction creates problems in all sorts of contexts; there is nothing special about the problems it creates for the semantic descent account. Moreover, though there is no consensus about the right one, there are perfectly reasonable strategies for dealing with fiction. Still, while I don't think there is a pressing need, a few further words about how one strategy could work might help to defuse the worry about unicorns.

From my point of view, the most promising strategy in respect of fiction is the one calling on the notion of pretence. Familiarly, this strategy allows that, in using a sentence with a fictional name or predicate, we do not really refer to fictional entities, but rather pretend that we are referring to real ones. A huge amount has been said about this strategy, and no doubt even this summary description of it could be faulted. But, even without looking too closely at details, its appeal in the context of my account of metaphor is obvious. Were a helpful friend to describe an adversary you find intimidating this way: *Walter is the hound of the Baskervilles*, it would be perfectly reasonable to find this metaphor intelligible by semantically descending to the poor dog who in the story is got up to be frightening, but is in fact harmless, and seeing this creature as qualifying *Walter*. Even granted that, harmless or not, there happens to be no hound of the Baskervilles, there is no special problem in pretending or imagining that he does exist, and therefore in taking this pretend individual to characterize *Walter*.

There is obviously a lot more to be said about the scope of pretence in claims like that about *Walter*; the use of fictional entities in metaphors will almost certainly complicate any pretence account of fiction. But I hope that what I have said is enough to convince you that fictional objects, when we pretend them to be real, can play pretty much the same role in semantic descent and qualification as any other object. Nor should this be surprising, given that we need to exercise imagination in semantic descent even when the relevant predicates have non-empty extensions.⁹⁰

⁹⁰ My comments about fiction have been brief, but they are important. As the last sentence of this section suggests, even the determinacy of objects got by descent from non-empty predicates calls on our imaginative abilities, our tendency to engage in fictions. Indeed, one way to think of metaphors—something that is encouraged by the semantic descent account—is as small-scale attempts at fiction or, perhaps better, narrative. As I have already noted, this means that one should not get too bogged down in the metaphysics of semantic descent. Creators of metaphor can be understood as telling us very short stories about objects, and their use in characterizing the world. These stories ask us, sometimes but not always, to imagine determinate objects, but they offer no guarantees about their existence, and my account is one about how these objects are used, not about whether we can bump into them.

3.4. More on Semantic Descent (and Some Notation)

The notion of semantic descent was introduced by reference to the well-known Quinean idea of semantic ascent. In this section, I shall say something further about this, introducing along the way some notation that will prove useful.

In a typical case of semantic ascent, one takes an ordinary subject-predicate sentence: (15) The sky is blue, and treats that very sentence as a subject of comment via the truth predicate, namely: (16) 'The sky is blue' is true. This counts as *ascent* simply because (16) is metalinguistic, but it counts as *semantic* ascent because, in Quine's words: 'The truth predicate is a reminder that, despite a technical ascent to talk of sentences, our eye is on the world' (Quine 1970: 12).

Quine did not envisage the kind of case I have been describing, but nothing he says requires all cases of semantic ascent/descent to have the feature he associates with 'true', namely that the higher level sentence says the same thing as the lower-level sentence.⁹¹ One could regard preservation of content as a special case, appropriate only to ascent/descent using the truth predicate. Clearly, in the cases of semantic descent that I have described, sameness of saying is not preserved: taken literally, Romeo's (R), and the claim about Tolstoy in (T), are certainly not equivalent to their semantically descended counterparts. That said, there is still every reason to count my suggestion as involving genuine *semantic* descent. Our encounter with the words 'sun' and 'infant', together with our recognition of their context as metaphorical, leads us to employ the relevant non-linguistic objects in new, 'hybrid' predicates which, for convenience, could be displayed using the following notation: (Rm) Juliet is the ↓sun↓, (Tm) Tolstoy is an ↓infant↓. Note: like quotation used in semantic ascent, the '↓' marker of descent works in pairs, but there are differences and these will be discussed.

The result in each case is a linguistic structure in which one element, the predicational part, aims to give information about the other element, the subject part. (Of course, in each case, the predicational part is a hybrid, consisting of some words and an object. But that doesn't affect the present point.) Even though the predicates in (Rm) and (Tm) are not the same predicates as those in (R) and (T), and do not therefore convey the same information, they share a subject matter with their undescended versions. Thus (Rm) like (R) is about Juliet, and (Tm) like (T) is about Tolstoy. The descent in each case is made with our eyes still firmly on the

⁹¹ Nor is it completely uncontroversial that such sameness of saying is preserved in semantic ascent using the truth predicate.

world—indeed on the same bit of the world—so Quine's condition for the semanticity of ascent/descent is still fulfilled.⁹²

The down-arrow notation displays semantic descent in the various examples, but more must be said about the rationale behind this notation before I can justifiably put it to further use.

When someone utters a sentence in a particular context, something about the words used, or features of the utterer's performance, or the context itself, or some combination of these, alerts us to the fact that the utterance involves a metaphor. Exactly *how* we manage this identification, as I have noted several times, seems to me an often messy empirical matter, best decided by (perhaps) psycholinguists, and I have concentrated instead on the philosophical task of giving an account of *what* is identified. Given this division of labour between the empirical and philosophical, the down-arrow notation should be seen as a tool for the philosophical theorist. It is intended to record—from the perspective of my particular theoretical account—what it is to be metaphorical, and should not be taken as a marker whose independent identification triggers the judgement that an utterance is metaphorical.

This warning is necessary, not least because down-arrows were compared earlier to quotation marks, and the latter are of course not merely creatures of theory.⁹³ But, so long as we keep this warning firmly in mind, the comparison between down-arrows and quotation is suggestive of something important. Quotation marks are a linguistic device that allows us to talk about the bit of language they enclose. If we remove them we get another linguistic device that allows us to talk about some appropriate bit of the world. Extending the process by using down-arrows around a predicative expression, we create yet another device, but one which now allows us to use one bit of the world to talk about some other bit.

Someone might hesitate to follow me here, insisting that 'semantic' and 'predicative' apply only to practices that involve words. The thought would be: when down-arrows are used, the result is an object, not words; that is, not something that we can use to *talk* about the world.

More or less this very issue featured prominently in the previous chapter: predication, being understood as essentially word-involving, I had to coin 'qualification' as the wordless undertaking of the predicational function. In so doing, I had in mind the idea that qualification should be counted the broader of the two notions, and that predication was therefore a species of qualification. The model here is reference: we manage to refer with objects, gestures, sounds, thought constituents, as well as words, and we don't even have a special term for reference by words. In contrast—and, as I argued, unreasonably—we think of predication as essentially something

⁹² In the more usual case of ascent/descent via the truth predicate, both directions are possible. Yet it can seem from what I have so far said about the metaphorical case that only descent has a point. This is not always so, as will be discussed in the section on dead metaphor in Ch. 4.

⁹³ Though given the widespread misuse of quotation marks, it may be simply optimistic to regard them as part of everyday practice.

achieved by using words, and we neither have a term for the more general activity, nor, partly for this reason, do we give it its due.

On the picture of qualification as the superordinate notion, there should be no problem in regarding the objects got by descent (in metaphors) as fulfilling a semantic function. In some cases, we use words to qualify items in the world (i.e. we use predicates); in others, descending from words to objects, we use these objects without the help of words to qualify items in the world. On the assumption that predicate-style qualification is a way of talking about something—a way of engaging in a semantical enterprise—there should be no problem in regarding object-style qualification in this same way.

That said, I could imagine someone finding this terminological proposal a little too radical. It is one thing to regard qualification as something like predication—close enough to find the parallel interesting—and another to treat qualification as the superordinate term of the pair. But all that I need in the present context, indeed, all I need for my account of metaphor, is the minimal concession that qualification is something like predication. Instead, therefore, of treating qualification as the superordinate term, we could think of it merely as a sort of precursor to fully semantic predication, that is, as ‘proto-semantic’ or ‘proto-predicative’. Not only does this not take anything substantial away from my account of metaphor, this way of putting the matter has useful resonances. After all, in so far as metaphor calls on a process which gives us insight into our ordinary notion of predication, it becomes itself implicated in the foundations of language. (I have already mentioned the foundational significance of metaphor, and it will figure more prominently in Chapter 4. At an appropriate point later in this chapter, I will return to the issue of the semanticity of hybrid or proto-predicates, since I certainly do not think that the remarks in this paragraph settle the issue.)

The idea of semantic descent, the proto-predicates it creates, and the down-arrow notation for displaying these predicates, are none of them everyday items, either in or outside philosophy. But there is a sort of language-like construction in common use whose familiarity might help here, *at least expositively*. Think about what is going on in the ubiquitous: I ♡_. For all practical purposes, the heart-shape in this structure has now become synonymous with the word ‘loves’, so I cannot claim that it offers us an actual case of semantic descent. However, one could imagine a time when there was no conventional link between the symbol and the word. Thinking back to that imagined time, we can imagine further the heart-shape being invoked as a device for referring to the heart, an object conceived of as, roughly, the seat of our affections. In my terms, one could think of the heart-shape as combining the functions of the down-arrows and any words they enclose. Instead of ↓sun↓, we have ♡_: the first picks out the object in the sky, the second the organ in our breasts. With this background, one could then finally imagine someone constructing the structure shown above. In this structure, the heart-shape picks out an object which itself then functions as a hybrid

or proto-predicate, one which qualifies the relationship between the speaker and whatever fills the blank.

Note that the heart-shape has a fixed effect, or one might say ‘meaning’: in every context, it picks out the heart. In contrast, we can think of the down-arrow notation as a functional device for turning almost any ordinary predicate into a proto-predicate. It offers us a way of concocting any number of symbols whose effects, in a given context, are something like that of the heart-shape. Rather than having recourse to the commercial artist's ability to think of expressive shapes, the least talented amongst us can turn words into such ‘shapes’. The down-arrow notation in effect marks the capacity of metaphor to harness the expressive power of language for uses other than that of ordinary predication. Also, the heart-shape example hints at something that I shall explore in the next chapter: serving as a two-place proto-predicate, a verb, it offers the merest hint that the semantic descent account, unlike many others, can cope with metaphors of widely differing syntactic forms.

3.5. The Route from Here

In the previous chapter, I introduced and defended the idea that objects can sometimes take on roles typically thought of as linguistic. In this one, I have set myself the task of using that idea in the semantic descent account of metaphor. What this account requires is, first, a movement from the words used in metaphors to objects, and, second, the use of these objects as qualificational, perhaps proto-linguistic, devices.

So far in this chapter, I have concentrated on the first of these requirements, that is, on offering some defence of the idea that the word-to-object movement in metaphor is something available and plausible. Further defence will of course have to be given. What I haven't yet done is to consider the issues raised by the second aspect of the semantic descent account. That is, I have yet to consider how qualification works in specific metaphors. In particular, I have yet to say anything in detail about the kinds of constraints that might make it reasonable and informative to maintain, for example, that the sun qualifies Juliet.

Saying the right things about both linguistic and non-linguistic context will be crucial here, but this is no surprise: context is bound to be a large part of anyone's story about metaphor. For instance, almost every writer who uses the example takes into account the fact that Romeo's assertion occurs in a larger linguistic context, one which gives useful, even necessary, guidance to interpreters. Of course, how this guidance is brought to bear will depend on one's account of metaphor. In my account, linguistic context works in part by giving us some idea of the kind of syntactical role an object is to play, and in part by hinting at its qualificational role. The sun might well work in one way in (R), set as it is in Romeo's extended monologue about a love-object, Juliet, and in another way, when this same object is used metaphorically in a story about the warrior Achilles.⁹⁴

⁹⁴ So far from being a problem, the fact that the sun can play different roles in this way seems just that little bit more evidence of the potential of objects to be predicates. If words can be ambiguous, here is evidence that objects can be as well.

No less, and often more, important is non-linguistic context. As with linguistic context, the function of this kind of context crucially depends on the account of metaphor in question. Though more will be said, it should be obvious that, on my account, the *focus* of non-linguistic context is, in each case, the object of metaphor got by semantic descent. That there is this focus for the information we think of as contextual is no small part of the reason to look with favour on my account. By insisting that there is such a thing as an object of metaphor, and that the purpose of this object is in fact predicative, we give contextual information something specific to work on and with. Accounts which merely point us in the direction of context, and suggest that it somehow helps us understand what is being got at in a metaphor, fall down badly in this respect.

Relevant contextual knowledge of the objects of metaphor can take many forms: it might be a matter simply of commonsensical knowledge, called on by the setting of the metaphor; or it might be a matter of conventional or even, to use a fancy term, *cultural* knowledge. After all, the sun, along with many of the other objects we call on in metaphor, can be described as having cultural significance (or significances)—a way of putting the matter which suits my purposes, given the resonance amongst *significance*, *sense*, and the wonderfully dextrous *meaning*.

Later in this chapter, I shall consider these matters further, though the nature and relevance of context to my account can only be fully appreciated when more complex examples of metaphor are in play in Chapter 4. However, aside from filling in details of semantic descent and qualification, there is a further important bit of unfinished business. In Chapter 1, I argued that each of the most prominent accounts of metaphor in the literature fail to accommodate at least one of the truths I described there as non-negotiable. It is thus crucial to show that the semantic descent account does better. For aside from worries about this or that aspect of the account, unless it manages the three truths about metaphor, it is no advance on those currently in the field.

It might seem best to take on the tasks in the order described above: first fill in some more details, and then show how the semantic descent account deals with all of the truths canvassed in Chapter 1. However, I have decided that the reverse order is actually better. Enough has been said about semantic descent (in this chapter), and the practice of qualification that grounds it (in this and the previous one), to make my overall account of metaphor clear, even if it needs further elucidation to be fully convincing. And it seems to me the best hope of convincing anyone is, first, to consider how an account shaped like mine negotiates the three basic truths, and then, once it is clear enough that some such account has a real future, to fill in the further details. Nor is this simply a tactic: some of the further detailed questions about semantic descent and qualification will be sharpened, even answered, by the discussion of the non-negotiable truths, and this will make the final sections of this chapter more rewarding.

What follows, then, in the next three sections will be a consideration of how my account manages to accommodate the truths about metaphor, followed in another section by further material on semantic descent and qualification.

3.6. Semantic Descent and Truth

The point about assertion and truth is easily accommodated. So long as one is careful not to equivocate, semantic descent can explain why it is perfectly in order to question, contradict, or assent to metaphorical utterances, and nothing in this explanation requires us either to eschew genuine metaphorical content, or to locate the focus of these truth-evaluating activities in something other than that content. Insensitive though it may be, an interlocutor may counter Romeo's claim, without changing the subject, by telling him: No, Juliet is not the sun. Unlike Davidson, I do not regard this remark as obviously true. For, as intended by the interlocutor, we must see this counter-claim as using the proto-predicate demanded by semantic descent, namely, 'is the ↓sun↓'; treating the word 'sun' here as part of an ordinary linguistic predicate would be to equivocate, thereby turning the interlocutor's remark into a bad joke.

The fact that the semantic descent account preserves the intuition about metaphorical truth should be unsurprising, since, as has been noted, it comes under the head 'Content Sufficient'. It belongs with the proposals of Beardsley, Black, Goodman, and Kittay (among others) who count the creators of metaphors as often straightforwardly aiming to communicate something appropriate to the words they use, and as responsible for the truth of what they assert. When Romeo says that Juliet is the sun, or when the critic says that Tolstoy is an infant, they are each using language to make assertions whose content makes essential use of our knowledge of the meanings of the expressions used, and which might be true, or might be false.⁹⁵ This applies equally to high-octane 'poetic' metaphors, even though special factors might well intervene to make the mundane, give-and-take practices of assertion inappropriate. The point about truth is not that we always insist on it in metaphorical utterance, but that we must find something to apply it to in those cases where it does figure. In this regard, it should be noted that even unarguably literal claims can be made in poetic contexts without our bothering over their truth or falsity. A poet who writes about certain flowers blooming in May is not thought to be speaking as a botanist.

Given its occurrence in a play, Romeo's remark in fact does occur in a truth-irrelevant context, but, for the purposes of the example, I have been following the usual convention of treating the remark as if made by a non-fictional character keen on conveying a truth about the object of his affections. Standing below the balcony, Romeo could have said something non-metaphorical about Juliet and, by that convention, this would have been taken as an assertion. (Imagine that he said: 'I love that woman.')

So, bracketing the theatrical context, there is no reason, intrinsic to

⁹⁵ Speech act views such as Searle's also see Romeo as making an assertion, but remember that I have classified indirect speech act views as Content Insufficient because content in the relevant sense is not by itself sufficient to explain what Romeo is doing. However, understood correctly, the content provided by the semantic descent account is sufficient for that purpose.

the metaphor, why we should regard (R) as anything other than an assertion for whose truth Romeo takes responsibility.

What makes Romeo's remark truth-apt is his qualifying use of the hybrid predicate 'is the ↓sun↓', and though this begins with a call on linguistic knowledge, knowledge of the copula, and the descriptive phrase 'the sun', the call on his and our knowledge goes beyond this. How far beyond? And in precisely what direction? These questions are in essence those I raised earlier about the processes of semantic descent and qualification, and, as noted, the plan is to deal with them later on. Still, I have to say something here about these questions, even if it is only to give the merest sketch, since I am aware that there might be a certain scepticism about the assertoric credentials of (R), understood in the way suggested.

The source of this scepticism is likely to be a certain model of linguistic understanding which, if left unchallenged, can make it seem as if there is a vast, even unbridgeable, difference between ordinary assertions and metaphorical assertions, at least as the latter figure in my account. Since it would take more than a chapter of another book to deal with this thoroughly, my devoting only part of a section to it in this one suggests rightly only the merest outline. But that outline is necessary, and I hope to make it clear enough for one to imagine how to fill it out.

Consider this perfectly ordinary assertion that Romeo might have made: (17) Juliet is a woman. What would we count as showing that someone understood an utterance of (17)? It is all too easy to think that in understanding what is said about Juliet here, what is required (at least in part) is a grasp of something—a meaning—which determines, among other things, how to sort items into those which fall under the concept *woman*, and those which do not. Or, since talk of meanings is out of favour, that what is required is that an interpreter bring to the context of utterance his or her knowledge of the contribution to truth conditions of the predicate expression 'woman', where this contribution is thought of as something which is in principle capable of dividing the world into those things which do, and those which do not, satisfy this predicate expression.

Of course, these can be perfectly innocuous claims: they may be taken as merely convoluted ways of insisting that anyone who understands (17) must know that it asserts of Juliet that she is a woman. But if one isn't careful they can be taken as the preface to something thoroughly misleading. I have in mind here the model of understanding which takes too seriously the idea of there being such a thing as the meaning of the predicate (perhaps its extension-determining power), and imagines it as a device which an appropriately trained speaker has somehow stored up in his mind, and is able to deploy when required. The trouble with this has nothing to do with the idea that such a device might be an abstract or mental object, and everything to do with the job of work assigned to it. For the idea of a meaning as a device which somehow contains the principle of sorting that goes with 'woman' is one that has been rightly criticized by Wittgenstein (and many others).

Of course, it is controversial that there is a problem with this story; many still find it worth telling, even given the Wittgensteinian rule-following arguments. (As far as I can see, psychologists tell this sort of story all the time, not even noticing that the Wittgenstein who gave them family resemblance to play with only a few pages later offered trenchant criticism of a use that has come to be made of family resemblance in the project of representing concepts.) In any case, my aim here is not to detail the arguments against the view—as noted, that is not a task for this book—but simply to warn against allowing it to influence, perhaps subliminally, your understanding of my account of metaphor assertion. For if you think that the predicate in (17) can only be understood by someone who has got hold of some such thing as its meaning or sense, and if you think of this meaning as some sort of device for determining the application of the predicate, then you are apt to be particularly unhappy about my account of Romeo's metaphorical assertion. You are apt to point out that there is nothing in the use of the object—the sun—which corresponds to such a meaning, nothing which fixes a range of application of this object when it is embedded in the linguistic framework of predication. Moreover, in not finding anything that answers to the meaning of qualifying object, you are apt to question anyone's taking Romeo's (R), interpreted in my way, as an assertion.

Several examples will show what I am up against. First, imagine a loyal retainer to the Capulet household who overhears Romeo's utterance of (R) and says: It is early and Juliet has just come out on the balcony. Romeo says that she is the sun, but he is deceived. Most mornings Juliet sleeps in until nearly noon: she is after all a teenager. Clearly, it is extremely tempting to describe such a case this way: the retainer takes Romeo to be asserting a thought, but it is not the one we imagined him as expressing. Hence the retainer's assessment of falsity is irrelevant, and communication non-existent. If in this case the retainer had made a simple mistake—perhaps just mishearing Romeo's words—then there would be no problem. But the worry hanging over the semantic descent account of metaphor is the possibility that this kind of thing could be the *norm* in cases where qualification figures, and hence that the qualifying object simply does not properly fix the extension of the hybrid predicate. Here is a second example. In correspondence, Jerry Fodor (helpfully?) suggested this version of Romeo's (R): Juliet is a real knockout; hot stuff, and, while this might somehow seem more on track than 'early riser', it too is unlikely as a rendition of what Romeo said.⁹⁶

What lies behind these examples is something like this. In the ordinary case of (17), we can describe an interpreter as having worked out that: (18) Romeo said that Juliet is a woman,

⁹⁶ It is also, at least faintly, metaphorical, but I shall let that pass here.

because this credits the interpreter with the possession of, among other things, the meaning of ‘woman’, and we can rely on anyone who, as it were, owns this meaning to be able correctly to sort things, in virtue of this meaning, into those which satisfy the predicate and those which do not. However, when it comes to (R), my recommendation is that we describe an interpreter as having worked out that: (19) Romeo said that Juliet is the \downarrow sun \downarrow , but this doesn't seem to come with the same guarantees as (18). That is, (19) does not credit the interpreter with possessing a meaning associated with the sun which would rule out the retainer's and Fodor's unacceptable renditions. Indeed, the worry is that (19), for all that I have said so far, could be used to characterize *both* the retainer's and Fodor's versions, even though these are plainly enough conflicting:

Responding to this, I suggest that the appeal-to-meanings gloss of the straightforward (18), the sentence we use to report an interpretation of (17), is misleading, and this for broadly Wittgenstein's reasons. However, when we understand it correctly, the way is also open to seeing (19) in a more flattering light.

When (18) is asserted by some interpreter, what it reports is not the possession of some extension-determining element called the sense or meaning of the predicate ‘woman’. Rather, it reports the fact that, by our lights, the interpreter has made sense, in a specific way, of the speaker's action in producing just those sounds—that the interpreter has managed to fit those sounds into a larger network of attitudes and actions. Here is a way of spelling this out: The adequacy of the total theory [of sense or meaning] would turn on its acceptably imposing descriptions, reporting behaviour as performance of speech acts of specified kinds with specified contents, on a range of potential actions—those that would constitute speech in the language—describable, antecedently, only as so much patterned emission of noise. For that systematic imposing of descriptions to be acceptable, it would have to be the case that speakers' performances of the actions thus ascribed to them were, for the most part, intelligible under those descriptions, in the light of propositional attitudes; their possession of which, in turn, would have to be intelligible, in the light of their behaviour—including, of course, their linguistic behaviour—and their environment. The point of the notion of sense—what the content-specifying component of a total theory of that sort would be a theory of—is thus tied to our interest in understanding—fathoming—people. We have not properly made sense of forms of words in a language if we have not, thereby, got some way towards making sense of its speakers. If there is a pun here, it is an illuminating one. (McDowell 1998: 172)

Going perhaps a bit further than this suggestion, but I believe in the same direction, I would describe (18) as a marker of a kind of co-ordination that exists within a linguistic group. The group consists of the speaker of (17), the interpreter of (17) who produces (18), and we who underwrite the interpreter's (18) as a correct interpretation of (17). In effect, our preparedness to accept (18) as true indicates our confidence that Romeo's action in uttering (17) does in fact fit intelligibly within the overall complex of *his, the interpreter's*, and our *own* actions and thoughts. We can say that the acceptability of (18) shows there to be a kind of *attunement* amongst all the participants. Note though that it is one thing to consider (18) as announcing or

marking such attunement, and another to insist that it, or elements in it, bring that attunement about. However tempting it is, (18) should not be thought of as introducing an element, for example, the sense of 'woman', whose possession somehow *dictates* co-ordination amongst the relevant parties. Though there is nothing about Wittgensteinian exegesis that is beyond dispute, I think that this is not only an important element in his rule-following considerations, but is one that could and should be taken on board by philosophers of language. In any case, it is something that I shall accept here, leaving further arguments in its favour for another time and place.

If we look at (19), that is: (19) Romeo said that Juliet is the ↓sun↓, in a parallel way, the problem raised by the retainer's and Fodor's versions turns out to be more apparent than real. For if we take (19) as announcing attunement amongst Romeo, an interpreter, and ourselves, rather than as having the role of imposing such attunement, nothing prevents our insisting that, as far as the retainer and Fodor are concerned, (19) is just not warranted. The logic of the original objection went like this: the retainer and Fodor were clearly mistaken about Romeo, but, since each could be characterized by (19), this attribution is just too thin to be an account of what Romeo said. To which I reply: given the proper way to understand such attributions, and the fact that we understand straight-off just how mistaken both the retainer and Fodor are about Romeo's (R), nothing compels us to employ (19); if there is no attunement, then there is no grounds for asserting what is, after all, by our lights a marker of attunement.

I can imagine someone thinking that this reply shows a certain perversity. Surely, the idea behind the objection is that the retainer and Fodor can *both* sign up to (19), because they both do in fact recognize that what is being said is that the sun qualifies Juliet. They do see that what is involved is a metaphor, and they cannot be accused of merely taking the original (R) to be a kind of literal, coded way of saying either that Juliet gets up early or is hot stuff. But, having signed up to (19), each of them goes on to make comments about Romeo's assertion that show them to have misunderstood. And the fault lies, so to speak, not with them but with my account of what Romeo said.

There are two things that can be said in reply here. First, I think that this way of putting the objection presumes just the kind of demand on interpretative attributions that I laboured to discredit. (It also shows how easy it is to slip into the mistake of finding such a demand reasonable.) It is simply not the case that when we find an interpreter able to assess a certain utterance as an assertion with the content, say, that a is an F, we are thereby crediting that interpreter with possessing a device which itself correctly determines the application of F. The right picture is really quite the reverse. It is because, by our lights and against a background of intelligibility-conferring attributions, the interpreter has got hold of a certain way of treating things as F, and shares this with the original speaker, that we find correct the interpreter's assessment of the speaker's utterance. Given what is in any case presumed by the objector

to my account, that both the retainer and Fodor do *not* share an understanding of (R) with Romeo, we should have no hesitation in resisting using our (19) in characterizing their understanding of his assertion. Insisting that the retainer and Fodor both know that (R) is a metaphor and even perhaps that it involves the use of a hybrid predicate ('is the↓sun↓') to qualify Juliet, is not enough to make (19) appropriate. It would be enough if the hybrid predicate was a device which itself fixed this application, so that anyone who, as it were, owned it—who knew that it was operative—couldn't make the retainer's or Fodor's mistakes. But this is precisely the picture that we should learn to ignore.

The second point to make in reply is somewhat more concessive. I do realize the worry about my account of Romeo's assertion turns in part on a difference between ordinary linguistic predicates and hybrid or proto-predicates. With the ordinary ones, for all that they keep philosophers of language up at night, we use them pretty unreflectively in making assessments of understanding. Romeo utters the sentence:(17) Juliet is a woman,someone hears this and says:(18) Romeo said that Juliet is a woman,and we have little hesitation in thinking that the interpreter got it right. The connections between, for example, Romeo's and the interpreter's linguistic actions using the word 'woman', and the no less important connections to the myriad further actions and attitudes of all those who use 'woman', are simply and succinctly crystalized in the deployment of the word itself. In contrast, there is no such easy route to the characterizatoin of Romeo's claim in (R). However, even while conceding this, I insist that this difference is no more than one of degree.

Think of how it can go wrong, even in the purely linguistic case. I have in mind here things the interpreter might go on to say which could make us wonder whether he and Romeo should be seen as occupying the same place in, as it is useful to put it, the space of reasons that the use of 'woman' marks. (Using Sellars's metaphor-laden terminology helps to reinforce the earlier point about attributions: Romeo and the interpreter do not come to occupy a location in the space of reasons fixed for them by their use of certain words. Rather, the words have the significance they do in virtue of speaker's and hearer's occupation of some such place.⁹⁷) For example, if the interpreter of the ordinary (17) went on to claim that Romeo was thereby saying:(20) Juliet is wilful,or:(21) Juliet is subservient,

⁹⁷ I do realize that, in using Sellars's phrase so soon after citing McDowell, and going so far as to mention Wittgenstein, I am connecting up my discussion to a literature that is becoming ever more vast, thus rendering my few remarks about this nexus of philosophical theses superficial at best. That said, I make no apologies for making the connection, since I think it important to the overall project in this book.

we might begin to wonder whether, as it is often put, he had the same concept of *woman* as Romeo. Of course, we can wonder this without having yet to give up on (18) as a correct interpretation. However, this is only the beginning, and there are other ways of going on that might leave us less sure about the matter. Suppose that this same interpreter insisted: (22) Romeo said that Juliet was the offspring of his uncle. This might make us wonder whether, from our point of view, the retainer had somehow confounded ‘woman’ and ‘cousin’. And more bizarre continuations can be imagined—continuations which would make it impossible to maintain the pretence that ‘woman’, as it figures in (18), is the right word to use in characterizing the interpreter's report of Romeo's utterance.

With my treatment of Romeo's (R), these same kinds of problem are perhaps more easily conjured up; suggestions like those of the retainer and Fodor do not require much imagination. But what is important is that they present problems no different in kind from those we can come up with by applying our imagination to the linguistic case.

On my account of (R), there is no way to capture the content of Romeo's assertion by using purely linguistic predicates. The linguistic means we employ to succinctly characterize the place in the space of reasons which, for example, Romeo in uttering (17), and his interpreter in uttering (18), occupy, just don't work for (R). Nonetheless, this should not lead us to abandon the hybrid-predicate characterization of that content. The predicate expression ‘woman’ contributes to the content of (17) because it is treated as reserving a particular place in the space of reasons; it functions, or is rather allowed to function, without the need for further commentary.⁹⁸ The hybrid ‘is the ↓sun↓’ does not get this same treatment—that was conceded above—but, when we look a little deeper, we can see it as no less fit for purpose. The very fact that we can see straight off that the retainer's and Fodor's suggestions are wrongheaded is evidence of this.

There are things we know (or believe) about the sun, as well as about the extra-linguistic and linguistic context in which Romeo's assertion was made, which should be seen as background to our use of the hybrid predicate in (19). These are precisely the sorts of thing that I will say more about in the last section of this chapter, but it is important here to have the right view of them. I am not saying that the things we know about the sun etc. are a substitute for things said to be known by language users in respect of ordinary predicate expressions like ‘woman’. Having resisted the picture of understanding ordinary predicates as the possession of devices (concepts?) which impose their correct uses on speakers, I am scarcely looking to fill in

⁹⁸ Which is not to say that there is something behind the use of ‘woman’ in (17) and (18) which guarantees that the speaker, interpreter, and ourselves are always going to be in step. The use of the word in these formulas indicates, rather than determines, co-ordination or attunement, so things could go badly wrong, even though they have gone swimmingly up to the point of our acquiescing in this particular interpretative use of ‘woman’.

a similar picture in respect of hybrid predicates. The point about those things we know about the sun, and about the background to Romeo's utterance, is not that they constitute such device, but that they help us identify the place in the space of reasons appropriate to the hybrid predicate as used by, and of, Romeo when he utters (R).⁹⁹ For all that we put our faith in the use of 'woman' in certifying the attribution in (18), that faith is grounded on the background knowledge we have of the things people think about the world, about each other, and of course it is no less dependent on the things they do, and the reasons for which they do them. Similarly, our faith in the hybrid predicate's potential to reflect the content of Romeo's (R) is grounded in such knowledge, notwithstanding a tendency to think that purely linguistic predicates can do their interpretative work without explicitly calling on any such outside help. Indeed, it would be surprising if there were not some such difference between literal assertions like (17) and metaphors. For if there weren't such a difference, metaphor would simply not be the phenomenon it is. (I shall say more about this in the next section, where I consider the issue of paraphrase.)

More promissory notes have been issued, but the upshot of the section has been a defence of our using down-arrow representations in specifying the truth-relevant content of the assertions made with metaphors. These representations reflect the background against which the qualifying use of the appropriate object is made.

⁹⁹ At the risk of introducing just that bit too much of McDowell, let me try to make the point about attributions, using his notion of 'sideways-on' vs direct (i.e. not sideways-on) perspectives. Everything up to now has been mainly about our view of the interpreter's view of Romeo's utterance, all of it therefore sideways-on. Think now about how Romeo might conceive of his utterance (apologies to Shakespeare). He sees the sun on a wonderful spring morning, and finds it to convey information about Juliet; the sun, he reckons (in my terminology) qualifies Juliet. In doing so, he is guided by what he knows about the sun, what he believes, what he believes others believe or know, etc. All of this leads him to make the assertion (R). In doing so, he expects to be understood. Why? Because (as my account has it) his words are straightforward and he expects that his audience will perfectly well understand the kind of background information relevant to this particular qualification. Now imagine the retainer and Fodor approach him: the first says, 'So you mean she gets up early', and the second 'So you mean she's hot stuff'. Here I expect Romeo to be exasperated with the retainer and Fodor, not with himself. From his point of view, he has a perfectly good thought, it was the reason he said what he did, and he would expect that others who share the background information about the sun, and understand him would also understand his utterance. Yet these two characters just didn't get it. Could he have somehow spoken sloppily? Did his words have an ambiguity he should have avoided? How could he put his uncomprehending auditors right? These are all questions from Romeo's point of view—a view that is definitely not sideways-on—but since it is my example, I will take the liberty of imagining some answers from my sideways-on perspective. Romeo did not speak sloppily or ambiguously, no more so than anyone does who uses words that require some sensitivity to what is going on. It was simply that the auditors did not show the requisite sensitivity. Surely, Romeo had a right to think that his auditors knew as much about the sun as he did, and also knew that he was talking about Juliet, the object of his love and devotion. There is a range of things his auditors could have said which might have surprised (even delighted him) because one can often find that one's meaning is best elucidated and expanded upon by someone else. But Romeo is right to think that the retainer and Fodor simply didn't get it. How to help them? Well, he could say some things about the sun by way of elucidation, he could say some things about Juliet by way of providing reasons for his remark. But as for the remark itself, nothing needs to be added: it is not the predicate to which the sun contributed that needs modification, just the auditors. (More on the contrast between the *elucidation* of a metaphor and a *reason* for it in the next section.)

They do not by themselves impose uses of the necessarily hybrid predicate that figures in the representation of the metaphor. But that is not something we should expect, even of purely linguistic predicates as used in specifying the content of common-or-garden literal assertions. Whether we are speaking of purely linguistic predicates or my hybrid ones, these figure, not as ways of imposing certain patterns of use, but as marking an attunement in thought and action between speaker, hearer, and all other participants in the space of reasons.

3.7. Paraphrase

According to my account, understanding a metaphor consists in fastening onto an object-exemplar of a relevant predicate, and then taking that object as a qualifier of the metaphor's subject.¹⁰⁰ In this way, a hearer who is attuned to the predicative use of that object comes to understand something about the subject of the metaphor, while justifiably considering the speaker responsible for the truth of that information. In sum, all the ingredients of common-or-garden assertion are present in metaphor. However, since the information is conveyed using a hybrid predicate—one which contains an object in a qualifying role—it would be bizarre to ask someone to express this same information *in other words*. Admittedly, it is words in the metaphor that call on the object. But it is what I have also called the 'proto-predicate', object included, which conveys a message, not the words themselves. Since the speaker is using an object, not words, to convey a message, it makes no sense even to try to paraphrase a metaphor in the strict sense of the term.

A useful way to think of it is that, on the semantic descent account, a metaphor functions like a picture, diagram, or map. As we saw in Chapter 1, a request to paraphrase a picture makes no sense, but then again neither does a request to paraphrase an object. Nonetheless, there is plenty of house-room in the semantic descent account for the other activities with which paraphrase is all too easily confused. I am referring here to translating, elucidating, and generally commenting on metaphors.

3.7.1. Paraphrase and translation

As already noted, it is not usually all that difficult to translate a metaphor into another language. This somewhat mysterious fact has been noted by many writers, but has never been to my mind satisfactorily explained. While native English speakers find it hard to imagine how, in general, the power and beauty of Shakespeare's language can survive translation, we do not have the same difficulty with certain bits of that language, bits such as Romeo's description of Juliet. Whether she is 'le soleil', 'il sole', or 'die Sonne' makes little difference to the impact of the original metaphor. Nor, on the semantic descent account, would we expect it to. For, on that account, what is being translated is the word or phrase

¹⁰⁰ Just to remind you: I don't think that metaphors are even mostly subject-predicate in form. For the present, however, I am not questioning this all-too-common assumption, because it makes exposition easier. In the next chapter, it will be shown how the semantic descent account copes with the syntactical variety of realistically complex examples.

before, as it were, descent to the proto-predicate takes place. So long as these translations preserve reference we should expect, rather than be surprised by, their adequacy.

One must be careful here: a superficial reading of my claim about preservation of reference can make it seem vulnerable to obvious counterexamples. For example, assuming that the sun is my favourite heavenly body, does this mean that we can translate Romeo's assertion into French as: (R1) Juliet est le corps céleste favori de SG? Or, given that the sun undergoes nuclear fusion, does this make the following a good translation: (R2) Juliet est un four nucléaire, autour de qui la terre fait sa révolution? Clearly, the answer to both of these questions is 'no', but then neither is really a serious counterexample to my original claim. What I contend is that, in cases where the words in a metaphor in a source language are replaced by their standard or usual translational counterparts in a target language, the semantic descent account can explain the surprising fact that metaphorical effect is preserved. The standard translational counterpart of 'the sun' in French is neither 'le corps céleste favori de SG', nor 'un four nucléaire, autour de qui la terre fait sa révolution'. So, though reference is preserved when these expressions are substituted for 'the sun', they are not reference-preserving *translations*.

Couldn't there be cases in which the only available way to render 'the sun' preserved reference, but led to hopeless translations of the metaphor? Suppose, for example, that in some language, Native, the best you could do for 'sun' would be something which came back into English as: 'the evil staring eye of the Ox-god who rises from bed every morning'. (For whatever reason, the speakers in this community regard the sun as threatening and malevolent.) Clearly, this: (R3) Juliet is the evil staring eye of the Ox-god who rises from bed every morning, disfigures Romeo's remark. But doesn't it also undermine my claim about translation? Someone might insist that we have in (R3) a translation which preserves reference but doesn't cope with Romeo's metaphor.

Though the matter is somewhat intricate, I think this possibility is no real objection to what I have claimed. My defence proceeds on two fronts. The first of these is straightforward: I have not insisted that *every* putative translation of an expression preserves the metaphorical content of the original; merely that when there is preservation, it can be most naturally explained by the semantic descent account. The second defensive front is less straightforward: there is good reason to wonder whether we should accept that, in respect of 'sun', translation between English and Native is genuinely possible. And, even if we do accept that translation is possible, the result is more supportive than undermining of semantic descent. More detail on both strategies follows.

When a hack, English–French translator, for example, one who makes a living translating washing-machine manuals, gets to work on a text of Shakespeare's, we would expect the result to be less than lovely. Yet, even with such low expectations, we would find that some bits of the text, for example, Romeo's metaphor, pass through the translation process with little or no loss. This is surprising enough to call for comment, and what I have been claiming is that the semantic descent account can explain, without fuss, why this happens *when it does*. I should not be taken as having made an open-ended claim about the possibility of translating *every* metaphor into *any* language. So much then for the first line of defence.

Is it appropriate to regard (R3) as a translation? The example was under-described, so it needs filling out. One possibility is that Native speakers have no word for the sun, always using a *phrase* in their language which, when retranslated back into English, leads us to (R3). In this case, I think we would very much doubt that English–Native translation of 'sun' is possible. The need for a circumlocution in one language replacing a single word in another is usually a good indicator of untranslatability, and if (R3) is not a translation in any reasonable sense, then there is no case for my account to answer. (That this could happen with 'sun' is highly unlikely, but there is nothing incoherent in making the assumption that it could happen.)

The other possibility is that, in addition to the 'Ox-god' descriptive phrase, there is a single word for the sun in Native, and hence that translation would seem possible. Strictly speaking, this would mean that something had gone wrong with the retranslation of (R) back into English as (R3): if there were a single word in Native for sun, the retranslation should come back as just (R). Yet, as the case was described, it doesn't seem quite acceptable to say nothing more than that retranslation gives us back (R). Clearly, the case deserves a closer look.

From the perspective of the semantic descent account, the one-word translation of (R) into Native lets us down, not because anything has gone wrong linguistically, but because the attitudes of Native speakers towards the sun, the object itself, differ so strikingly from our own. This is shown by their readiness to employ a phrase that comes back to us as in (R3). That this can happen is familiar enough: there are non-hypothetical cases in which translation is unhelpful, in spite of referential overlap in vocabulary. French and English have words for cabbage, but the French use this vegetable to characterize human beings in a complimentary way, while this is most certainly not something done by English speakers. If something like the vegetable scenario were true of Native and English in respect of the sun, then it would certainly be unreasonable to deny that translation is possible. But it would be no less unreasonable if we thought of these translations as preserving metaphor.

Does admitting this undermine the idea that semantic descent explains the relative ease of translation of metaphor? The answer is 'no', partly because of my first line of defence: semantic descent explains why it is often easy to translate metaphors, but this does not imply that translation always works, let alone that it is always easy. But there is more to it. For not only does the semantic descent account explain translation when it is easy, it also gives what seems the right explanation for

failure in cases where translation doesn't work. And this is something that many other accounts get wrong.

Black, and many Content Sufficient theorists who follow him, try to account for metaphors like (R) by complicating the story we tell about the meanings of the words. Very roughly, Black invites us to think that words not only have minimal and narrowly fixed (literal) truth conditions, but that language users are able to construct further meanings for these words on a case-by-case basis as these words occur in metaphors. The crucial thing to note is that the focus of all this extra effort is *words*: it is the familiar 'sun' that comes to have an unfamiliar meaning in (R). Thus, when someone attempts to translate (R)—which, from our point of view contains 'sun' with its newly minted meaning—into Native and then back into English, that result is bound to be unsatisfactory, given what the natives think of the sun. For this reason, a Black-style theorist will come to think that translation isn't possible, even if Native has a single word for the sun. There is thus no room in a Black-style account for the obvious point made earlier that, if Native has a word for the sun, the problem is not translation of 'sun', but rather the attitudes that Native speakers have towards the *object* that 'sun' picks out. The semantic descent account makes objects the focus of metaphor, whereas for most other Content Sufficient accounts this focus is words, a difference that is actually quite important. (Something like this same issue will figure in Chapter 5 when I discuss Josef Stern's account of metaphor.)

3.7.2. Paraphrase and other near relations

Moving on from translation, I turn now to something merely touched on in the treatment of metaphorical truth. Metaphors—especially but not only of the literary variety—tend to provoke explanatory commentaries, and it is crucial to understand how these commentaries relate to what I have been calling 'paraphrase in the strict sense'. For unless I clarify this, there will always be the suspicion that I am perversely hanging onto a much too strict idea of what is involved in paraphrase.

Here are some of the things that might be produced in respect of Romeo's remark:(23) i. Romeo thinks that Juliet is necessary to his very existence.ii. Romeo thinks that Juliet is responsible for his seeing the world aright.iii. Romeo thinks Juliet is time itself.(24) i. The sun is the ultimate source of light and warmth.ii. The sun is the measure of time.iii. The sun makes life on earth possible.Though I have perhaps been overzealous in displaying these comments as belonging to sharply bifurcated classes, there can be little doubt that some such division exists. Some comments sound like glosses on what a speaker/writer aimed to achieve in using the metaphor, whereas the others take aim at the metaphor itself. Neither should be seen as a way of paraphrasing what the speaker said, though it is often difficult to resist this when the comments take the form of (23i–iii). My suggestion

is that we can make best sense of (23)–(24), and their difference from paraphrase, from the perspective of semantic descent.

As I have insisted, everything begins with our taking seriously the idea that Romeo made an assertion, that he expressed a thought he regards as true. On the semantic descent account, this expression of thought, achieved as it is through metaphor, is special: the thought expressed is available to a hearer who is attuned to the metaphor, but we should not think that this requires a hearer to be able to find other words that express that same thought. This is because Romeo's thought is conveyed by a proto-predicate—an object pressed into service as a qualifier of Juliet. Quite simply: though Romeo used familiar words *in* expressing his thought, the thought itself was not expressed *by* those familiar words.

That said, Romeo's assertion is undeniably an intentional action, something undertaken for a reason. It can be seen as aiming both at expressing something about Juliet and (in principle) at conveying that something to an audience. (One must put on one side here the fact that Romeo is, as it were, talking to himself. The complexities wrought by the fact that this particular assertion occurs as part of a monologue in a play are interesting but not relevant to the general point being made with this example.) Now, as with any other action, it always makes sense to wonder why it was done—to ask after the state of mind of its agent. Moreover, given that we are here dealing with a *linguistic* act which, unlike, say, making yourself a cup of tea, is expressive and informative, it makes sense to wonder about the effects that it can have on an audience.¹⁰¹

Putting these considerations together, it is unsurprising that commentaries take the form they do. The examples in (23) are typical of the attempts one might make to understand the state of mind which led Romeo to say what he did, whereas the examples in (24) suggest things about the metaphor object which might be the basis of attunement between speaker and audience. However, in neither case should we see these claims as aiming to capture the thought that Romeo expressed; in neither case is the aim to put Romeo's assertion into *other words*. Still, it is easy to mistake these attempts for paraphrase, especially when it comes to the items listed in (23). Such a mistake might take the form of someone's insisting that I have left out an important category of 'commentary'. The objector continues: in addition to (23) and (24), you should have included:(25) i. Romeo said that Juliet is necessary to his very existence.ii. Romeo said that Juliet is responsible for his seeing the world aright.iii. Romeo said that Juliet is time itself.This is because at least part of the reason Romeo has for uttering the sentence: 'Juliet is the sun', is that he thinks that Juliet is necessary to his very existence. So, we should count this thought as part of what he actually expresses when he does

¹⁰¹ Given the central thesis of this book, I am the last person to insist that non-linguistic acts are not expressive and informative. So, take the example in the spirit intended: making a cup of tea might well be expressive, but not always.

produce the metaphor. By doing this, we can achieve paraphrase, albeit partial and approximate, by uttering the sentences in (25).

My first reaction to this—one that I expect is not idiosyncratic—is that there is something distinctly odd about the sentences in (25). It is certainly right to think that Romeo asserted something, but it doesn't seem right to think that he actually said any of these things. There are of course complex (and disputed) constraints on accurately reporting speech, but they are surely not met by any of (25). To sharpen this perception, suppose you hadn't heard Romeo's remark, but were told: He said that Juliet was necessary to his existence. What range of words would you then imagine Romeo to have used? Not the ones he actually used, I suspect. Of course, my finding (25) odd might well be an artefact of my advocacy of the semantic descent account. So, it is important to say something more.

The sentences in (23) are a sample of the kind of thing that Romeo might well have thought, and knowing these thoughts certainly helps us understand why he said what he did in uttering (R). However, there is no reason to think that thoughts which in this way *explain* an assertion are themselves *expressed* in it. To be sure, this mistake is easily made: a reason for my saying: 'The meeting was short but tedious' is that I thought it was. Here the thought behind my utterance, and the thought expressed in it, are one and the same. But, as should be obvious enough, this is by no means generally true. For example, thoughts that explain an assertion might in some way imply, or be implied by, though not be the same as the thought expressed in it. Or, they might be thoughts about some end, the envisaged achieving of which brings about the expression of some other thought, perhaps one characterizing the means. My suggestion is that metaphor is more plausibly counted as in the class of assertions whose supporting reasons are distinct from the thought expressed. Moreover, this suggestion is not dependent on acceptance of my account. Everyone agrees that the thoughts expressed in, or by, the words of a metaphor are in some way distinct from thoughts expressible without metaphor. Indeed, the perennial mystery and appeal of metaphor is in large part traceable to this elusiveness. If for no other reason we should therefore be ready to find sentences such as those in (25) off-beam. That Juliet is necessary to Romeo's existence may well be a thought that helps us gain a deeper insight into Romeo's remark and/or one which is part of his reason for making it, but this wholly non-metaphorical claim can scarcely be a whole or even a part of what Romeo actually asserted.¹⁰²

The temptation to think that perfectly apt commentaries on metaphor are also ways of paraphrasing them is very difficult to resist. After all, the aim of commentary

¹⁰² If, try as hard as you can, you cannot come up with anything *else* asserted by Romeo, you are in particular danger of taking sentences like (25) as genuine attempts to put Romeo's thought into words. This is what makes the semantic descent account particularly important in the present context: it gives us a way to preserve the irreducibility of metaphor, while still making space for elucidatory and explanatory comments about particular metaphors.

is a deepening of understanding of both a metaphor and its author. Without care, it would be easy to slide from ‘understanding’ in this broad sense to ‘understanding’ as, more narrowly, grasping what is said. The upshot of this would be that commentary becomes paraphrase. But if one holds fast to the idea that a metaphor is in some way a device for communicating thoughts distinct from any purely word-based one, then there is plenty of scope for allowing commentary that is in no way paraphrase.

Just so that we can keep track, I shall call comments like those in (23) *rationalizations* of a metaphor, and those in (24) *elucidations*. I do not insist that this classification is exhaustive, though it is certainly generous enough to allow in most of the kinds of thing you are likely to come across in discussions of any specific metaphor. In the present context, the thing to note is the way in which the semantic descent account not only shows us how to keep these two kinds of activity separate from paraphrase, but also how to make the distinction between rationalizations and elucidations.

Rationalizations and elucidations can take various forms and, while this is not the place to sift through actual commentaries, sorting their contents into these two categories, we can count the sentences in the somewhat artificial lists (23) and (24) as pointing out the way to do it. What is suggested by these two lists is that: (i) rationalizations are essentially non-metaphorical attributions of thoughts about the subject of the metaphor, either to the speaker (as shown in 23), or to an audience (not shown, but easily imagined); whereas (ii) elucidations focus, not so much on this traffic of thought, but on features of the metaphor that make such traffic possible.¹⁰³

While one can easily imagine commentaries that combine these two features, they are nevertheless distinct, and the semantic descent account provides a straightforward explanation of this. Given the idea of metaphor as a device for reaching down through language into the world, and using what is found there to express and convey information, we should actually expect to find commentary taking (at least) these two forms. On the one hand, there would be rationalizations—considerations focusing primarily on what was called above the ‘traffic of thought’. Here what is in question is not so much the content of the metaphor assertion as the trail of thoughts that leads to and from it. On the other hand, there would be elucidations, comments which take aim at the metaphor itself and, as I see it, highlight the roles of the relevant objects in what are proto-predications.

Unlike the case of ordinary predication, in proto-predication one cannot rely on attunement marked by what are presumed to be shared words. In the case of

¹⁰³ It might be helpful to contrast the two notions in a context somewhat removed from metaphor. One could think of rationalizations as like the comments that a film director makes (typically on a separate track on the DVD recording) explaining what he hoped to achieve with a particular scene. Elucidations are more like the comments that a film critic might make about how that particular scene achieved whatever effects it did. Note though that there is absolutely no temptation to count either of these as a ‘paraphrases’ of the scene and, though the context is different, the explanation of this is pretty much the same as in the case of metaphor, namely, the fact that the scene itself did not say anything in words.

metaphor, attunement must be grounded on a shared appreciation of the significance of objects. Our attunement to Romeo's remark comes not from what we know about the meanings of words—though words and their context do play an important part in the process—but from what we know about the sun. In so far as this attunement is less than perfect, or perhaps simply to make it more explicit, one would find it natural to make the kind of comments represented in (24). Indeed, it is difficult to make sense of these sorts of comment except on the semantic descent account. (My comments on attunement have so far been rather teasingly vague; a more focused discussion of this will round out this chapter.)

Summing up this section: it has been argued that the demand for paraphrase in respect of metaphor is simply inappropriate. Admittedly, this is because I have been insisting on taking 'paraphrase' in a strict sense, but I have offered grounds for this. Coming at the matter indirectly, I have argued that the only real pressure for loosening our standards of paraphrase comes from two interrelated sources: (i) an unreasonably narrow model of assertoric communication, one which requires every thought expressed in such communication to be expressible, even if approximately, by words in natural language; and (ii) a natural, and therefore intelligible, but ultimately wrongheaded tendency to take the kinds of thing we say about specific metaphors as themselves paraphrase (though in a 'loose' sense).¹⁰⁴ Taking these sources into account, and treating them with the suspicion they deserve, we can relieve the pressure on the notion of paraphrase. Neither rationalization nor elucidation of a metaphor is any more a way of paraphrasing than is translation. Moreover—and this is the most important bit—the semantic descent account helps us understand what is going on in respect of all of these notions. It suggests why translation of metaphor succeeds, against the odds for translation generally; it grounds what is a real distinction between rationalization and elucidation within the category of metaphor commentary; and finally it offers a straightforward reason to stick with the notion of paraphrase as it is ordinarily and strictly understood.

3.8. Transparency

I begin with a brief résumé of my discussion of this feature of metaphor, as it is the least familiar of the three, and hence most subject to misunderstanding. Also, some of the points made above in connection with truth and paraphrase now put us in a better position to understand transparency.

When you hear the utterance of a sentence in your native language, then, assuming nothing untoward about background noise or trouble with your ears, you

¹⁰⁴ Though I haven't mentioned this so far, and don't want to make too much of it even now, it really does seem odd to speak of 'loosening' our standards of paraphrase (in connection with metaphor) when the notion itself is defined in a way which simply doesn't allow such talk. As noted, any dictionary will tell you, and as is obvious from the word itself, paraphrasing is saying the same thing in other words. While it is clearly reasonable to think that the 'sameness' referred to can be tighter or looser, dropping the implication of the phrase 'in other words' is not so much loosening the notion as abandoning it. And, given my remarks about commentaries, this is not something we need to do.

just do hear the meaning of the utterance. This is not something voluntary or optional: it is just not possible to suppress the meaning of what is heard.

This familiar fact is the beginning of research in psycholinguistics—research aimed at explaining why this is so and how it is accomplished—but the fact itself is something familiar, and not the upshot of such research.¹⁰⁵ ‘Transparency’ is the label that I use for this feature, albeit with some misgivings, and what I insisted in Chapter 1 was *both* that transparency is as true of metaphor as it is of unarguably literal sentences *and* that many accounts of metaphor could not cope with this parity.

In showing why certain accounts of metaphor cannot cope with transparency, the claim itself needed some refinement. This did not involve resiling from my insistence on transparency as the involuntary hearing of meaning, but it proved necessary to exercise care in explaining the import of the phrase ‘hearing meaning’, in order to avoid either trivialization or implausibility.

There are two threats on the trivialization front. On the one hand, it is not enough to say that the meaning we take in on hearing appropriate utterances is simply literal. It would be convenient if this literality claim were true. When Romeo utters: Juliet is a woman, English speakers cannot avoid hearing this as having the content that Juliet is a woman. If in the same way, when Romeo utters the familiar: (R) Juliet is the sun, the content taken in is that Juliet is the sun, that is, that she is the very large far away object in the centre of the solar system, we would be home and dry as far as transparency was concerned. Admittedly, what would be heard as a meaning in this case would be something absurdly false, but it would at least be a meaning or content. And this might be felt to be enough to explain the parity in respect of transparency between the literal and the metaphorical.

Unfortunately, this just doesn't work. Aside from the fact that there is a real issue as to whether the literal reading of (R) is nonsensical, not simply false—whether, that is, there is such a reading—there is an even larger problem. While I do not want to claim that, in hearing the meaning of (R), an ordinary auditor gets right to the bottom of this utterance, it certainly does seem true that more is got out of it than the merely literal. This is shown by the fact that whatever it is we take in on hearing (R) is enough to make intelligible the kinds of things said by way of what were

¹⁰⁵ Rather surprisingly, Fodor suggested (in an email) that my point about transparency was ‘straight empirical’ and therefore inappropriate as a philosophical objection to other accounts of metaphor. The surprise (to me) is simply that he more than anyone else seems willing to count empirical and therefore contingent theses within philosophical argument. But, more to the present point, he was mistaken in thinking that transparency is any more ‘empirical’ than the fact that metaphors are often taken as assertions, or that we find it odd to be asked to paraphrase a metaphor. For an account of some of the psycholinguistic research on metaphor see Gibbs 1994.

described in the previous section as elucidation and rationalization. I won't here repeat the arguments offered Chapter 1 to support this claim, but their upshot was that there must be some, perhaps shallow, kind of genuine understanding which preserves parity in respect of transparency between the literal and the metaphorical.

What this shallow understanding comes to is not easy to describe, and this brings me to the 'other hand' in respect of the trivialization danger. For, as we also saw in Chapter 1, it is tempting to try to characterize this understanding by saying that auditors hear the meaning of the familiar words in (R), and also understand that they are to take (at least some of) these words 'metaphorically'. But this is no improvement on the previous attempt. The problem is that, in finding there to be further work that auditors must do to get at the content—even the shallow content—of (R), transparency is lost. (Again, more detail on this was given in Chapter 1.)

Against the background of the failure of these simplest proposals for dealing with transparency—proposals that would make it trivially easy to satisfy—one has to be careful not to credit hearers with access to a content so profound as to be implausible. As has been mentioned, meaning is non-optionally heard in the utterance of (R), but it is most certainly not the kind of content that makes elucidatory or rationalizing commentary on the metaphor otiose. What transparency requires of accounts of metaphor is that they come up with a notion of content relevant to the metaphorical aspirations of utterances, but which nonetheless leaves plenty of further room for commentary on the nature of these aspirations.

It was at this point in Chapter 1 that I worked up a proposal based on Davidsonian materials to show that the transparency challenge could be met, a proposal I called the 'Image Account' (IA). Aside from its intrinsic interest, the IA has a special role to play in the present context, for I shall argue that there are overlaps in respect of transparency between the IA and my own account of metaphor.

3.8.1. Transparency: the image account and semantic descent

According to the IA, Romeo's (R) is not heard as some absurdly false literal claim, nor as a merely skeletal arrangement of familiar words accompanied by general advice to seek a metaphorical interpretation. In fact, it is not heard as a claim at all. Instead, we hear Romeo's utterance as: (R') Juliet as the sun. What (R') does is to put in front of us a way of thinking about Juliet, something we can call an 'image', so long as we do not focus solely on the idea of visual images. There are two features of (R') which make it ideal for meeting the transparency challenge. First, (R') suggests a depth of understanding that is just right for transparency. The image of Juliet as the sun connects to the kinds of elucidation and rationalization we are likely to make about this metaphor, but doesn't, implausibly, pre-empt them. Indeed, the IA allows for the possibility that someone could 'get' (R), but be stuck for anything else sensible to say about it. Second, there are good reasons to think that someone might actually hear (R) as (R').

As noted in Chapter 1, when we hear (meaningfully) a literal sentence such as: Henry is English, it is not implausible to think that there is something like a simultaneous activation of an image (again, not by any means a visual one) of Henry *as* English—an activation which, depending on the context, may be reinforced or extinguished. There is, after all, only a small lexical difference between the two sentences, and their contents resonate with each other. If you are told that Henry is English, this conjures up the image of him as English, and if you are offered only the image of him as English, it is difficult not to wonder whether he is English, even if this information hasn't been explicitly provided.

Of course, for all the talk of simultaneous 'activation', I do not intend to prejudge results of as yet unperformed psycholinguistic experiments. But there are experiments suggesting that we do in fact simultaneously 'access' multiple meanings of ambiguous words, and only later use context to disambiguate them. So, it is not that far-fetched to think we might do something similar with 'Henry is English'.

As noted above, my original purpose in constructing the IA was to show that transparency could be accommodated, not to offer the IA as genuine rival to my own account of metaphor. After all, the IA fails to negotiate the truth aptness of many metaphors, and, as described above, it is limited to metaphors of subject-predicate syntactical form. However, the usefulness of the IA is not limited to showing *that* transparency can be accommodated. In accomplishing this feat, it happens also to provide materials that the semantic descent account can appropriate.

The IA sees the words of a metaphor as, so to speak, projecting an image of one thing as another, and inviting a comparison between the properties of the objects that figure in that image. In contrast, the semantic descent account sees these words as forcing us to confront an object—an object which then functions as a proto-predicate (i.e. qualificationally), conveying information to an attuned hearer. This difference aside, however, the IA and semantic descent share a crucial feature: they require for their respective projection and confrontation an initial grasp of the ordinary meanings of the words used, and both take this grasp to be all that is required. It is this that allows them, in their otherwise different ways, to accommodate transparency. Relying as they do only on the ordinary meanings of the words in a metaphor makes it plausible that there is something available right from the start which can confer intelligibility, and this is all that is needed for transparency.

That said, the differences between the accounts do matter. The IA requires a shift from the copula 'is' to the preposition 'as'. This move invites hearers to explore, rather than to agree or disagree with, what is put before them. Of course, having made this shift, hearers can then use the words in the sentence with their perfectly ordinary senses: it is after all the properties of the objects picked out by these words that serve as the starting points of exploration. But, passing over the question of truth aptness, the shift itself, small as it can be made to seem, is more radical than that required by semantic descent. In the latter, the structure of the original sentence is preserved.

While this way of looking at the matter seems to be to the advantage of the semantic descent account, one must be careful (and honest). As noted above, the shift required by the IA is the syntactic one involved in the move from the copula to the preposition; all the other words in the metaphor undergo no shift in meaning. In apparent contrast, while the semantic descent account does not require any global syntactical shift, and while it begins with the ordinary meanings of the words in the metaphor, it does end up envisaging what might be thought of as a meaning shift, not of the whole sentence but of certain of its words. Though it is the ordinary predicate expression ‘sun’ which figures in (R), it is the sun itself which does the work needed to make the metaphor intelligible. Does this apparently additional shift diminish the account's ability to handle transparency? Does the IA then come out better in this respect? As you would expect, the answers to these questions are negative.

Though my account lends itself to idea that, in (R), ‘sun’ is in effect ambiguous, the second sense of, for example, ‘sun’ in (R) is so close to the surface that understanding, at least in some immediate sense, should be effortless.¹⁰⁶ It is surely plausible that on hearing ‘sun’, we have unproblematic access to the object for which this word stands. (More on this follows below in 3.8.2.) And given that the copula in the original sentences invites us to take this object as a *monadic* proto-predicate, we have all we need for an initial understanding of the metaphor, one that satisfies the requirement of transparency.

The IA accommodates transparency by taking the words in a metaphor with their perfectly ordinary senses, not to express a thought, but to put an image before us. Once it is understood that some such image is involved in ‘getting’ the metaphor, a hearer has all that is necessary to engage in further elucidation and other kinds of comment. For all that the semantic descent account differs from the IA, it accommodates transparency in much the same way. Words are taken to have their ordinary senses, but rather than images, these words lead us to objects. ‘Getting’ a metaphor—hearing it as meaningful—then consists in recognizing that the objects fulfil a semantic function and, as with the IA, even this minimal understanding meshes smoothly with further elucidation and comment.¹⁰⁷

¹⁰⁶ While I have no objection to speaking of ‘ambiguity’, perhaps this is not the best term to use, since what we end up with is closer to polysemy, i.e. multiple *related* meanings. Or, because semantic descent countenances two kinds of meaning—that of the ordinary predicate-expression and that of the object as proto-predicative—perhaps one should speak of *bisemy*. Whatever the label, what is important is that semantic descent shows precisely how the meanings are related, and it thereby makes manifest and precise the active role of ordinary meaning in metaphor. This is in contrast to most appeals to polysemy in metaphor, where one finds only hand-waving in the direction of those relations.

¹⁰⁷ An additional weakness of the IA account is what it leaves unsaid. While I allowed it to be intelligible why relevant elucidations and commentaries fit with the deliverances of the IA, this is actually a rather grey area. Without further guidance, it is unclear what we are to do with the images which, according to the IA, confront us. In contrast, though someone might feel that objects do not pin down interpretations in quite the way that linguistic predicates do, it is at least clear what is required of them. (Not that I am suggesting here or in the text that hearers are aware of the semantic function of objects as such. As with the examples given in Ch. 2, we often take objects to be qualifiers without necessarily realizing or labelling what we are doing.)

3.8.2. Psycholinguistic speculations

As suggested above, the IA shift from ‘is’ to ‘as’ is one that might well be unreflectively accessed by a hearer. In saying this, I was mindful of psycholinguistic experiments which show that multiple interpretations can figure right from the start of sentence comprehension. Still, since the experimental results I had in mind concern *word* meaning, not the meaning of whole utterances, I was stretching a point. With the semantic descent account no such stretching is required, since all that need be in play are the two interpretations of the expression ‘sun’; one, the wholly linguistic predicate, the other, the object specified by that predicate, but serving itself proto-predicatively. Connected as they are, it would be strange if they didn’t resonate with each other from the beginning of any sentence comprehension task. Nor is this all that can be said about the differences between the IA and semantic descent that might show up in psycholinguistic experimentation.

As I have said several times, the time-course of comprehension of metaphor is not strictly relevant to transparency, since transparency is about how metaphors strike us, not about how long they take to do so. Hence, my asides about these matters should be viewed as no more than speculative. Still, it should be counted as a mark against any account of metaphor that, in accounting for transparency, it forces one’s hand about such matters before we even set out to do the experiments. Thus, it may be that metaphors on average take longer to comprehend than equivalently complex literal utterances, or it may be that there is no real difference in the time-course of the two tasks. Neither of these options should be closed off, yet such closing off seems inevitable on the IA.

The resonance of the two ‘meanings’ mandated by the IA—the literal and the image interpretations—is a resonance between contents of *whole* utterances. In order to understand Romeo as having presented an image of Juliet as the sun, we have first to confront the idea that she is the sun. That is, we have first to understand (R) in some sense other than that of image presentation. Given this, if the IA were correct, metaphor interpretation would be systematically slower than literal interpretation. This point needs careful handling.

Earlier I suggested that the image interpretation might become available *during* the comprehension of (R), and I still think this might be so. But even if this speculation were true, it would not eliminate the time-lag of metaphorical interpretation. Nor was that its point. Davidson, among many others, claims that we typically search for the metaphorical interpretation *after* we have found that the literal reading is unhelpfully false. If this were true, it would certainly require metaphor comprehension to lag systematically behind the literal, since the identification of metaphor would depend on the appreciation of falsehood, and the process of metaphor interpretation would then depend on that identification. It was to avoid the latter dependence that I outlined the possibility of a resonance between assertoric (literal) and image presentation contents associated with (R). With such resonance, there is no need for an auditor to hear the falsity of the literal as a starting gun for finding the metaphorical interpretation. Instead, factors having nothing to do with the falsity of literal assertion might reinforce the metaphorical interpretation

that accompanies the literal. But, though this might be a way to avoid the anyway implausible dependence of metaphor interpretation on recognition of literal falsehood, it does not deal fully with the time-lag issue.

Remember the psycholinguistically relevant issue is not the time taken over the literal and then the metaphorical comprehension of any given metaphor. Rather, it is the average time taken for comprehension of flatly literal sentences (of various syntactic complexities) in comparison with the time taken for sentences that are metaphorical (but match the literal ones in complexity). Clearly enough, if there is a dependence of metaphor comprehension on the discovery of the falsehood of the literal, then metaphors are always going to take, on average, longer to comprehend than merely literal sentences. But even without this dependence, the IA still requires some kind of whole-utterance interpretation, or even, merely, appreciation, as a step on the way to revealing the image that resonates with the literal, and this creates a built-in time-lag in comprehension.

The effect of this can be seen even in a literal sentence. Presented with ‘Henry is English’ we have to get some way to understanding it—though without necessarily evaluating its truth—in order to be in a position to switch our attention to its embedded, image-presentation form: ‘Henry as English’. It would be a speculation too far to say what effect this is likely to have on the actual time-course of comprehension, but the very fact that the whole sentence has to be processed in some way in order for the image to be appreciated makes it likely that the image interpretation would take longer to process. And if this is true of a literal sentence, it is bound also to be true of a metaphor. In contrast, the semantic descent account is not committed one way or the other to any difference in processing times.

On the one hand, it could be the case that we need to process the whole of Romeo's metaphor before we take on board its embedded proto-predicative use of the sun. It is even possible that, in spite of my scepticism, something like falsehood or anomaly are needed to trigger the move from the predicate ‘is the sun’ to ‘is the ↓sun↓’. In either case, we would then expect a systematic time-lag in the processing of metaphors as compared with literal sentences. On the other hand, the predicate expression ‘sun’ might well be taken in the way I suggested earlier, namely as having two related meanings that are both activated *before* whole-utterance comprehension. If this happens, then there would be no reason why metaphors would take systematically longer to comprehend. In this they would be like the two interpretations of ‘John found a bug in the room’, both of which are available, so to speak, at the same time, precisely because the activation of both senses of ‘bug’ takes place before any attempt is made to process the whole sentence.¹⁰⁸

There is a final point necessary to round out this section. Because the IA requires the shift from ‘is’ to ‘as’, there is no obvious way to adapt it to the demands of metaphors not of the subject-predicate form. While it might be possible to cook

¹⁰⁸ It is not clear how one could do the experimentation necessary to decide between these alternatives, but a look at certain forms of cross-modal priming (e.g. ‘sun’ priming the sun) *might* help us decide on the plausibility of the bisemous activation hypothesis.

something up for more complex cases, it wouldn't have the elegant simplicity of the original IA—the changing of a single word—and that simplicity is crucial to its success with transparency. As far as the semantic descent account is concerned, the prospects are brighter. So long as the account can negotiate the move from subject-predicate to more complex forms—the first order of business in the next chapter—there should be no special problem with transparency. This is because all that the semantic descent account needs for transparency is the kind of bisemousness shown by 'sun'. Whatever the syntactical position of the word or words that are the subject of descent, so long as there is the possibility of both readings being available from the start, transparency will be preserved.¹⁰⁹

3.9. Attunement

A key ingredient in the semantic descent account is the idea that objects can function in very much the same way as linguistic predicates. An obvious challenge to this proposal is that this use of objects is not sufficiently constrained to fix the content of any assertion that might be made with them. Already touched on, this challenge bears re-statement, as a preface to facing up to it more directly than I did in section 3.6.

When someone says: Juliet is a woman, the content of this assertion is said to be fixed (at least as far as the predicate is concerned) by the meaning of 'woman'. However, when Romeo says: (R) Juliet is the sun, then, in order so much as to countenance the idea that the sun can take on a predicative function, we must find something that fixes the contribution of the sun to this assertion, something which achieves what meaning does in the case of 'woman'. Unfortunately, the challenge pessimistically concludes, it is difficult to see how any such thing could be provided. The sun is, as one might say, a mere object, not something which possesses the semantic properties necessary to fixing the content of assertions in which it figures. In respect of (R), this is shown graphically by possibility that there could be auditors who hear this utterance perfectly well, and even recognize the role that I claim for the sun, while still getting completely the wrong end of the stick about what Romeo is saying.

In section 3.6, I insisted that this pessimistic assessment takes too literally the idea that linguistic predicates come *equipped* with meanings, roughly devices the

¹⁰⁹ As you will see, metaphors of whatever syntactical form still display something of a 'this-is-that' nature, so it is all too easy to imagine the IA working just as well in these more complex cases. However, this is a mistake. What is at issue in transparency is comprehension of a sentence with a specific form. That there is a useful way we could think about the content of that sentence if we ignore its actual form might well be true, but would be disastrous for transparency.

mastery of which co-ordinate interpretations amongst speakers and hearers. Instead, I suggested that we should think of meanings, not as bringing about such co-ordination, but as reflecting the co-ordination that happens to exist. When an interpreter says: Romeo said that Juliet is a woman, and Romeo agrees with him, and we agree with both of them, all of this is simply a way of displaying the fact that, in so far as 'woman' is concerned, our linguistic actions (and the reasons for them) are co-ordinated with one another. In the language I used in section 3.6, the use of 'woman' in these interpretative claims marks out a place in the space of reasons occupied by all of the relevant parties. Or, putting it in a less fancy way, what this shows is something that is as close to common sense as one gets in philosophy of language: speakers and hearers in fact use 'woman' in the same way.

The upshot of these considerations is that we should not feel disappointed, even cheated, by the fact that objects don't also come equipped with meanings which impose order on those who use these objects predicatively. Someone hearing Romeo's (R) might very well get the wrong end of the stick about Romeo. But we can tell this by what that individual goes on to say and do, and this is in the end not significantly different from the way it works in the case of linguistic predication. Someone using 'woman' might also get the wrong end of the stick, might, as is said of the Wittgensteinian pupil, go on in the wrong way, by using 'woman' in ways that surprise and confound us.

Even if you accept the above lesson, and come away thinking that there isn't a radical difference between linguistic predication and the use of objects to fulfil that same role, you might well feel that there is still difference enough to call for comment. I couldn't agree more, and it is the business of this section to face up to this fact. Interpreters use ordinary predicate expressions like 'is a woman' to display the co-ordination that exists in the use of that word, and I think an interpreter is perfectly justified in using 'is the ↓sun↓' to display a similar co-ordination in respect of Romeo's (R). But such co-ordination does not come about miraculously, and there are differences between these two cases.

No one would expect me to explain in a few pages, let alone a paragraph, how it comes to be true that users of linguistic predicates expect, and are right to expect, that their own use of these predicates matches the use made of those predicates by other members of their community. This is because what is necessary to filling in this picture is nothing short of an account of how human beings come to learn or acquire a natural language. Still, I can say something which is anodyne enough to be uncontroversial, but which nonetheless highlights an important difference between linguistic predication and non-linguistic qualification.

Given a certain genetic endowment—one which is probably language-specific, as Chomskyans maintain—we acquire the lexicon of this or that natural language by training in, and/or exposure to, uses of the items which make up that lexicon. More particularly, we all use 'woman' in much the same way because we are much the

same (genetically and in other ways) and we have had something like the same kind of training in or exposure to the use of this word. One of the most obvious differences between linguistic predication and qualification shows up here.

It has been my contention that objects, where this includes individual things, actions, events, states of affairs, etc., are used by human beings to convey information about other objects. In another place, I should like to consider at length whether the genetic endowment which allows us to use objects in this way is part of the same dedicated (or modular) endowment which makes natural language possible. It seems reasonable to think this true; after all, if there is some specialized genetic basis to our abilities to use *words* for reference and predication, it would be surprising if this didn't also underpin our ability to use non-word *objects* in these same ways. (Admittedly the referential use of objects is familiar, while their predicational use is not. But I hope that discussion and examples in Chapter 2 show that this disparity is superficial.) Still, whatever may be true of our genetic endowment, it seems obvious that, while we acquire predicate expressions in the lexicon through some kind of training, nothing quite like this happens in respect of objects. Even if I am right about the use of the sun in (R), and in other utterances, I wouldn't claim that we acquire this ability in precisely the same way we acquire the ability to employ the sound 'woman'. How important is this difference? Does it undermine the whole idea that objects can fulfil predicative functions? While it is easy to anticipate my replies to these questions, I will return to them, after considering a further potential difference between predicates and objects.

I suspect that many would regard expressions like 'woman' as relatively immune from the unruly demands of pragmatics. There is a kind of fiction that philosophers of language have often adopted according to which pragmatics makes its contribution, when it does, within certain circumscribed areas of the linguistic realm. Appeal to pragmatics is fine, indeed necessary, for fixing the reference of certain elements in utterances, for tense, and for coping with the potential unclarity that comes from the short cuts we take in speaking, for example, saying 'the table' when there are many such things in the immediate and further off environment. But, aside from the often forced examples of radical pragmaticists, the contribution of 'woman' to 'is a woman', and from thence to the truth conditions of sentences in which the latter predicate operates, does not seem to call on pragmatics.¹¹⁰

In Chapter 1, I thought it best to steer clear of debates about such matters as the semantics/pragmatics distinction (divide?), and I still think this best. But the point that I need to make here doesn't require any heavy-duty theorizing. It is simply this: in so far as you think a predicate expression like 'woman' is substantially, or even wholly, a matter of semantics and not pragmatics, you are bound to be disappointed by qualifying uses of the sun. The latter, as with most instances of qualification, are seriously context-dependent. On my account, we need to see the sun as informative

¹¹⁰ I have in mind examples discussed in Travis 1981.

in order to understand Romeo's utterance, but we must make liberal use of pragmatic considerations in doing so.¹¹¹

In sum, typically linguistic predicates tend to be acquired by some kind of training, and our dealing with them is rather more semantic than pragmatic, whereas no one thinks we have somehow been trained to use the sun in our understanding of (R), and the context-dependency of the latter is evident and unavoidable. These are real differences—and I will return to the issue of pragmatics and context below—but what I shall suggest now is that they are not as important as they might seem.¹¹² The underlying mechanisms of co-ordination in action and thought—co-ordination that that I called attunement—are really quite similar, whether that attunement is marked by words or objects.

3.9.1. Training and culture

On the simplest model, words are things we are trained to use, and because of this we find it natural to think of the co-ordination marked by words as largely the product of this training. This seems to contrast sharply with the case of objects; it can seem more a matter of luck than anything else when the use of an object as a predicate by one person is taken up in that same way by another. There are two strategies that one might use to soften this contrast.

On the one hand, one could question the role of a narrow and simplistic notion of training in language acquisition. On the other hand, one could insist that there is something very much like training, now understood more liberally, discernible in the many ways we think about objects. Since these strategies are obviously compatible, even reinforcing, I suggest that we adopt them both.

Chomsky argues forcibly that 'training', as this notion would ordinarily be understood, is neither obviously present, nor could be efficacious, in our acquisition of language. Wittgenstein ridicules the idea of training as a practice in which a teacher inculcates a naïve pupil's responses by giving explicit instruction. There is no doubt that Chomsky and Wittgenstein come at this issue from very different directions but, taking a wider perspective, one can see a certain sameness in their conclusions.¹¹³ What is crucial to both is a conception of what is brought *to* a putative training session. In their different ways, they subvert the simplistic idea of training by arguing that human beings bring to training what is in effect the very thing that training is supposed to inculcate. In particular, it is argued that the subject of

¹¹¹ One author (Stern 2000) goes so far as to posit a special operator to account for metaphor—an operator which converts ordinary predicates into context-dependent ones. While I shall discuss the problems of this view at some length in Ch. 5, and while I agree with him that context is crucial to metaphor, I can say now that I disagree about the point on which context exerts leverage.

¹¹² They may well account for what I claimed to be the mistaken idea that linguistic predicates, in contrast to my objects in a qualifying role, have meanings the mastery of which brings about attunement. However, these differences are trouble enough on their own, and I am trying to keep them separate from the myth of meanings that Wittgenstein and others have laboured to expose, and which I discussed in section 3.6.

¹¹³ See Guttenplan 1992 for an attempt at reconciling Chomsky and Wittgenstein, even in the face of Chomsky's resistance to any such thing.

linguistic training either possesses language already (in some sense), or comes equipped to respond, perhaps simply by being human, in the ways that the teacher demands.

Even without going into detail, it should be clear how this kind of subversion serves my purposes. Suppose that we use ‘woman’ as we do, not primarily as a result of training, but because we somehow are born with that concept, or, less problematically, that there is a place for some such concept already prepared in our cognitive economy. Then the difference between linguistic and non-linguistic predication will seem vanishingly small. Attunement in our thoughts about objects such as the sun is not something explicitly inculcated: it depends instead on our common humanity, on our reacting to the world in similar ways. But the considerations above suggest that something like an appeal to common humanity is no less necessary for understanding ‘woman’.

The admittedly highly general talk about common humanity introduces the second strand of my softening-up strategy. However, here even more than before, I am forced to make what are only schematic remarks. For the simple truth about a very complicated matter is that the kinds of objects my account calls on in metaphor are typically enmeshed in rich, and often acquired, patterns of thought and behaviour. Objects called on by metaphors are not merely particulars, actions, events, and states of affairs with which we causally interact; they are often bearers of what I earlier rather weakly called ‘cultural’ significance. In ways that are more the concern of anthropologists than philosophers, different societies have somehow managed to turn many of these objects into what are typically called symbols.¹¹⁴ In some cases, for example, that of the sun, the network of symbolic associations is fairly obvious. But in more interesting cases, it is more deeply buried. To take just one example, certain polarities such as that of night and day, or male and female, have had far-reaching, deep, and unobvious effects on the ways in which we classify things that are not themselves straightforwardly related to these (or other) polarities.¹¹⁵

This is no place for even the most potted account of cultural or symbolic significance, nor am I qualified to give one. In any case, spelling out examples of our susceptibility to these significances is not the issue. What is important is understanding how we come to be susceptible in the first place. For my suggestion is that, when one takes on board scepticism about simplistic ideas of training, the way is clear to seeing our acquisition of cultural or symbolic significances as not all that different from our acquisition of linguistic forms. To be sure, I am not denying that there are

¹¹⁴ I hope it is obvious that my account of metaphor can be understood as supplying a refinement to the often overly flexible term ‘symbol’. Very roughly, symbols are objects that take on semantic functions. Sometimes, a symbolic object is a simply referential device, but, if I am right, the really interesting case is that in which an object takes on the work of predication.

¹¹⁵ Here let me acknowledge an enormous debt to Mary Douglas. In her written work, she has made a substantial contribution to our understanding of what is sometimes called ‘analogic’ thinking, a topic which connects directly to my account of metaphor. But, further, I have been fortunate in having been able to have many discussions with her more directly about metaphor.

differences; I had better not be, if I am to remain on good terms with Chomsky and his thesis of language-specific genetic endowment. But there is no harm in pointing out that, as in the case of language, we almost certainly end up finding significances in objects because of a potent mixture of hard-wiring (i.e. genetic endowment), sharing of reactions (without explicit instruction) and training, in a more defensible sense, that can be fitted under the rubric 'culture'.

3.9.2. Significance and context

While still allowing there to be differences between linguistic predication and the use of objects as qualifiers, I have tried various strategies to minimize them, to treat any such difference more as a matter of degree than anything more radical. But for all that I have said, I don't think this is the end of the matter, and I need to confront several obvious further worries, not to say objections, which are in fact interconnected. Doing so will lead directly into some further remarks about context, the real business of this subsection.

As put to me on several occasions, one objection is as follows: language is compositional, productive, and, in these ways, systematic. Part of the point of developing a theory of meaning for a language—part of the point of speaking about a semantics for a language—is precisely to reflect these features. Can I really be saying that the meanings—call them 'significances'—of objects, when used as qualifiers, share anything like these features?

Consider these sentences: Juliet is a woman. George Sand is a woman. In each case, the predicate expression 'woman' makes the same contribution to the truth conditions of the whole sentence, and—as many have argued—the fact that they do make the same contribution is absolutely essential for explaining the compositionality and, hence, the productivity characteristic of natural language. In contrast, the objection continues, consider: Juliet is the sun, The warrior Achilles is the sun, J. M. W. Turner is the sun. According to the semantic descent account, the sun contributes subtly different things to the truth conditions of each of these sentences, differences that depend in part on the varying contexts in which each is asserted, and in part on the place or places of the sun in our culture. I have spoken of qualification as semantic, but is it really reasonable to think of the vast background network supporting each qualifying use of the sun as admitting the kind of systematicity necessary for compositionality?

The reply to this is not difficult. Of course, there can be no question of the significances of objects being organized in a compositional way. The context-sensitivity of

objects employed in qualification (where this includes their cultural roles) rules out any such thing. Moreover, if one insists on using ‘semantic’ so that the only things which exhibit this feature compose in the way that words in natural language are reputed to do, then I had better refrain from using this term of our ability to use objects as qualifiers of other objects. However, while these answers suggest something like surrender, this is not so. The lack of compositionality of objects in their qualifying roles is only problematic if it threatens whatever compositionality we think we find in natural language. But, on my account, no such threat exists. Indeed, it is precisely one of the virtues of that account that it leaves the words of natural language—meanings and all—unreformed. Black and theorists like him tell a story about how words, in the context of a metaphor, come to acquire new meanings. It is this kind of project which is threatened by the worry about compositionality, not mine. If ‘woman’, ‘infant’, or ‘sun’ figure in a compositional theory for English, this is well and good. The semantic descent account takes these words with their compositional meanings just as they are, using them only as conduits to the things these words, in their unreformed employment, extend.

As to the term ‘semantic’: while I could just allow that this term is appropriate only of natural language words (and constructions), and then only when they form a system recognizable as compositional, this seems unreasonable and contrary to common philosophical usage. Surely, if someone tells a story using a salt cellar to refer to their car, we would regard that as an exercise of a semantic ability, though we would certainly not wonder whether salt cellars compose. Similarly, if one sees some event as giving us information about another, then, however context-sensitive and unsystematic, we should I think regard what happens, just as we did in the referential case, as an exercise of a semantic ability.

Even if, in respect of acquisition and general functioning, there is no radical difference between words and objects functioning like words, and even if the issue of compositionality is not relevant, there is still a difference that needs further discussion.

One might be willing to accept that there are parallels between words and their meanings, and certain kinds of objects and their significances, but still insist that objects like the sun need pragmatic (contextual) input to avoid a kind of indeterminacy that simply doesn't attach to predicate expressions like ‘woman’. After all, it is one thing to fend off as mistaken interpretations of (R) which, for example, take Romeo to have said that Juliet gets up early, or that she is hot stuff. But the object itself still seems to leave what many would regard as an unacceptably wide range of less off-beam interpretations that can only be reined in by appeal to context. In contrast, though we can entertain the possibility of someone using ‘woman’ in a way which makes us doubt whether it is a woman being spoken about, this has nothing to do with an intrinsic indeterminacy in the use of this predicate expression—an indeterminacy that can only be eliminated by some kind of contextual constraint. Indeed, part of the point of this worry is that ‘woman’ simply doesn't have a point of attachment for context.

There are two ways to respond to this. The first accepts that there is this sort of indeterminacy in cases of qualification, and that qualification depends on

context in a way that ordinary predication doesn't. But it goes on to insist that, when we understand the nature of the contribution context makes, we should be less inclined to find the admitted difference between predication and qualification to be radical. (I have already made scattered remarks about context, but I will pull them together a bit more immediately below in spelling out this response.) The second response, less concessively, doubts whether this kind of indeterminacy is quite so intrinsic to the use of objects as qualifiers. The claim here is supported by my earlier insistence that there is something misleadingly simplistic about examples like (R) and (T), examples that have carried the burden of exposition in this chapter. (I will expand on this in the next subsection.)

Appeals to pragmatic factors to remove indeterminacy in interpretation can seem ad hoc. This is of course not so when the relevant word, phrase, or sentence is explicitly understood to need contextual supplement. Thus, pronouns, tense, demonstrative expressions, to take three well-known examples, have, as it were, a built-in requirement for contextual specification. Someone simply doesn't understand expressions like 'he', 'was', or 'that sofa' unless he recognizes them as context-dependent; where these items are concerned, no one would think appeal to context is ad hoc. Other words—predicate expressions like 'is a woman'—do not have what we can think of as 'hooks' for the attachment of context. When therefore someone gestures to context to deal with a problematic sentence containing one of these words or phrases, we would seem to have good reason to be suspicious.

I do realize that the distinction between expressions that have hooks for contexts and those that don't is apt to be controversial. However, it is not meant to be a way of reinstating the idea that the only kind of context worth considering is one that can be 'semanticized', one which demands explicit parameters onto which contextual information can attach. I can easily imagine there to be expressions demanding a contribution from context, even without this contribution being encoded in the explicit way required for demonstratives, tense, pronominal reference, etc. In any case, the distinction is only intended to be rough and ready. Still, even after these concessions are made, there does seem to be a real difference between the treatment of ordinary predicates, and those in metaphor. There can scarcely be a more obvious candidate for being a predicate expression that is *not* context-dependent than 'is the sun'; it is difficult to imagine there being hooks in this phrase that context could attach itself to. Even so, virtually all of the Content Sufficient accounts that I am aware of insist on a role for context in helping us understand the contribution this expression makes to (R). However, unless some story can be told about why we should regard context as relevant to understanding this apparently context-independent predicate, this insistence is little more than a gesture.

As mentioned earlier, there is one writer who recognizes this worry. Stern suggests that we must treat 'is the sun' in a metaphor like (R) as subject to an operator ('M-that') which turns this context-independent predicate into a context-dependent one. One could say that his operator furnishes the hook for context. A more detailed examination of Stern's view comes in Chapter 5, but it is worth re-emphasizing something here that I said in an earlier footnote. There is something suspect about

the idea of an operator that imposes context-dependency on a predicate expression that ordinarily gets by without it. If ever the charge of ad hocity was justified, this would be a case. Stern will get his day in court in respect of this charge, but what matters here is not Stern's view but my own. What I shall suggest is that the semantic descent account gives us a way to beat this particular rap: it shows not only why context is important to metaphor, but provides a principled way to see why it is important, a way which doesn't depend just on the fact that metaphors are unintelligible without it.

It should be noted that the above worry about Content Sufficient accounts is not the well-known one about whether it is right to complicate the meaning of predicate expressions in dealing with metaphor. Those who, like Davidson, cleave to an austere account of meaning, have objected to this kind of complexity and have suggested that we don't have a clear idea of how it can be accomplished. But if you reject austerity, and think that the key to metaphor lies in a more sophisticated notion of meaning, you won't be put off that easily. However, the problem with context raised above is not that it makes meaning theory complicated. It is that, while we know that a sophisticated account of meaning needs to take context into consideration, when it comes to predicates like 'is the sun' there is no obvious place for context to get a grip.

On the semantic descent account, this whole issue is neatly bypassed. My account is not in the business of generating new meanings for predicate expressions: 'old' (austere) meanings do just fine. Thus, in (R), 'is the (a)' marks monadic predication, and 'the sun' does just what one would expect: it takes us to the object, the sun. However, because it is the *object*, not the original predicate *expression*, that is charged with the information-conveying task in (R), there is nothing to prevent our recognizing, and much to encourage, the importance of context in the fulfilment of that task. Indeed, we can say that, in contrast to predicates like 'is the sun', objects in their information-conveying modes are thoroughly context-dependent, something true independently of any need to confer intelligibility on this or that metaphor. The mere fact that an object is called upon to serve as a qualifier is by itself reason to count that object as having a 'hook' onto which context fastens.

As always, when qualification is in question, it helps to think of cases in which objects take on the other role in the basic combination, namely reference. When someone tells the story of a recent accident using an item on the dinner table to stand for his car, he might begin: 'The salt cellar is my car ...' Now, the words used in this description are like 'the sun' in this sense: these words do not make any explicit demands on context for their intelligibility. To be sure, we need to look to the surrounding circumstances to tell which item is in fact referred to by 'the salt cellar', but we need to do this to be *au fait* with how things are, and not because we would otherwise fail to understand the description. However, while context is not necessary for understanding the descriptive phrase, it is most certainly necessary for understanding the key referential activity that is played out on the dinner table, namely that the salt cellar refers to a certain car. It would be odd to speak of the salt cellar itself as having hooks for the attachment of context, but it is not odd to say

that the act of using the object in this referential way comes with a built-in demand for contextual help.¹¹⁶

Details aside, the crucial moral of this referential case is that the mere use of an object as a referring device is obviously and naturally understood as dependent on context. There is nothing ad hoc about this dependency; it is not that we find ourselves insisting on the need for context as a result of finding unintelligible the communicative act in which the reference takes place. It is simply that any such use is a use-in-a-context.

This same moral applies to cases of qualification. Independently of metaphor, and of the demand that this or that metaphor be made intelligible, one recognizes the use of objects as qualifiers as a context-dependent activity. Hence, while ‘the sun’ is not obviously bound to any context for its interpretation, when used as a qualifier, the sun most certainly is.

With these remarks, one can better appreciate the real difference between my account's appeal to context and the appeal made by other Content Sufficient accounts. In the usual such account, and Black's can serve as an example here, contexts of the most diverse kinds are ransacked for ‘associated commonplaces’ that then help to make the metaphor intelligible. These commonplaces, if that is ever what they are, are those claimed to be associated with the predicate *expression* in the metaphor. But as we have seen, the predicate expression in a metaphor is typically not one that has any need of context for its intelligibility; context is thus sort of forced on it. In contrast, in my account, context is tied, not to an expression, but to the object that then comes to serve as a qualifier. Because qualification by objects is a function that is intrinsically context-dependent, there is no forcing of contextual ingredients onto an unwilling recipient.

Here is another way to think about these issues. Ordinary predicate expressions like ‘the sun’ do their predicative work (obviously, in this case, when combined with the copula) largely on their own: context is not necessary for them to function in their communicative or expressive settings. Focusing on this can make objects as qualifiers seem hopelessly indeterminate, because we simply do not think of an object as having a meaning matching that of a predicate expression in natural language. However, this is not the right way to look at qualification by objects. As we have seen, such qualification, like its referential counterpart, is intrinsically (and naturally) context-dependent. So, the right thing to compare with a linguistic predicate is not the object itself, but the object in its context. (Or, to be more precise, we should compare predicate expressions with the object in its contexts. The plural is required because, in some, but by no means all, cases, there can be systematically varying qualificational uses of one and the same object. This is no different from the

¹¹⁶ One might think then of the use of the salt cellar as a bit like a use of a demonstrative: both have built-in demands for contextual clarification. However, I put this remark in a footnote, partly because demonstratives raise many intricate issues that are not relevant here, and partly because I want to avoid claiming that object-reference involves the same kind of explicitness of appeal to context as demonstratives.

kind of multiplicity of meaning that attaches to natural language predicates such as ‘bug’.)

We should then not think of metaphor as requiring appeal to the object got by descent, the whole of which then requires appeal to context to confer intelligibility on the metaphor. This can make my proposal sound ad hoc in just the way that Black's is. Rather, we should think that intelligibility requires, from the beginning, an appeal to object + context. And, though this will vary with the specific example, many of the parameters of context which serve the use of an object as qualifier are in large part determined by the object itself; this is what it means to take seriously the idea of ‘object + context’. So far then from being ad hoc, we should think of objects as bringing to metaphors certain qualificational potentials. This reinforces a parallel between linguistic predicates and qualification that I have insisted on throughout.

While the above remarks made liberal use of ‘context’, there are distinctions to be made. Once it is recognized that appeals to some kind of context belong with, are required by, each qualifying use of objects, we should separate this kind of context from the more general kind that could figure in any use of a sentence, whether metaphorical or not. (I say something about this below.)

There are two main sources of contextual constraint that focus directly on the use of an object as qualifier, that is, which figure in the formula ‘object + context’. The first is largely linguistic, and the second largely a matter of shared, broadly empirical, knowledge, but this distinction can be blurred in specific cases. On the linguistic side, one begins with the fact that the words from which semantic descent is made exercise some control over the way the relevant object is used. We understand the sun to function monadically because ‘sun’ is embedded in the ‘is a (the)’ schema. (As you will see from the examples discussed in the next chapter, polyadic predication is no more problematic.) Additionally, the qualifying use of the sun is controlled by the fact that that very word is used for this object, instead of one of the many possible co-referential alternatives. The control in this particular example is subtle, and is best brought out by imagining ‘sun’ in (R) replaced by ‘astronomical body at centre of our solar system’, ‘nuclear fireball 193 million miles away’, or even ‘astronomical object worshipped by the Egyptians’. Each of these would transform and disfigure Romeo's (R), even though, in each case, the object got by semantic descent would remain the same. The plain fact, already discussed, is that the sun has a cultural significance for us, though we would be deflected from appealing to this significance if the object were referred to in one of these different ways.

The non-linguistic, but focused, context is made up of the knowledge (and of course beliefs) we have about the object which figures in the metaphor, though these are much too coarse as ways of describing the kind of underlying attitudes we have towards objects that I earlier described as ‘cultural significance’. After all, objects count for us in ways that we often don't notice, and would have difficulty in recovering without help from anthropological friends. Still, in so far as our differential appreciation of objects affects our behaviour and thought, and given the tendency to use ‘knowledge’ even when it is in some sense tacit, there is no harm in speaking of our knowledge of an object's cultural significance.

What I prefer to think of as ‘general’ context consists basically of the general factual and linguistic setting within which the metaphorical utterance is made. For example, Romeo doesn't just come out with (R), he says a lot more about the sun before and after the famous line. We also know, in his case, and do generally, lots of things about the speaker and the things spoken about. These include the fact that the subject of the sentence is Juliet, a young woman with whom Romeo is infatuated, that his emotions are in turmoil, etc. Many writers appeal to this kind of general context in connection with metaphor, but the simple truth is that such context is a part of utterance interpretation generally and, though in any given case it might be crucial, it has nothing special to do with metaphor. In fact, though I haven't gone in for the kind of detail necessary to force home this claim, I count it an advantage of the semantic descent account that it allows us to mark distinctions among different kinds of appeal to context. On the one hand, there are contextual ingredients which go with the object of metaphor and help to fix the contribution it makes. On the other, there are contextual ingredients which form the background to the whole of the utterance, and which, for example, help us to disambiguate ambiguous expressions (or objects). There is no doubt that the interplay between these two can make it seem as if there is just one kind of contribution that context makes. But this would be a mistake.

3.9.3. Dispensing with culture

I have used metaphor examples sparingly in this chapter. Romeo's claim about Juliet and the critic's swipe at Tolstoy have carried virtually all of the expository burden, and these two are syntactically almost as simple as metaphors get. While this minimalist strategy has the advantage of not distracting attention from the main lines of my account, it has drawbacks. Although I think we should recognize the role of culturally shared significance in respect of the sun or of infants, these examples are each rather special. It would be wrong to think that my account stands or falls on them, or indeed on the parallel I have suggested between the cultural significances of these particular objects and the notion of word sense. The sun is certainly an object which figures in all sorts of ways in our collective psyche, but that is just the trouble: it is such a flexible symbolic friend that its use in (R) needs heavy-duty support from all three kinds of context, it being one of those cases which give evidence of something like ambiguity. The situation is different with infants, but no more helpful: there can be little doubt that there is a shared stereotype of infantile behaviour, but we think so many other things about infants that context is crucial for triggering this stereotype. Additionally, the very fact that one can speak here of ‘stereotype’ suggests that (T) is conventional, and this might make one wonder whether, as already noted, the metaphor in (T) is dead. The subject of dead metaphor gets a thorough airing in Chapter 4, and I shall also have something to say in that chapter about the more interesting variant of (T) that the critic actually produced. But the point I want to make now is that there are endless examples of metaphor which can lend support to the semantic descent account, and in which context, especially the element of cultural significance, plays a smaller part. This is the less concessive strand of my defensive strategy.

Though many more interesting examples will figure in the next chapter, I shall close this one with a taster, something to serve as an encouragement to watch this space. Consider the following claim that might have been made by an observer in the Montague household: (27) Romeo is an elastic band stretched to its limit. Or, if you think this too anachronistic, try: (28) Romeo is an unsecured fifty-pounder on the starboard side with the ship set to come about onto the port tack.¹¹⁷ Unlike (R) and (T), these examples do not depend on being embedded in a rich linguistic context, nor does anything seriously anthropological get a look in. My account requires us to find the object of descent from the phrase ‘elastic band stretched to its limit’, and this is perfectly straightforward. It then requires us to imagine using this object as a qualifier of Romeo. Again, this is straightforward. Even linguistic control of this qualification—control usually exercised by the words from which descent is made—is minimal. We can tell from these words that the relevant object is intended to function like a one-place predicate, but, excuse the pun, there isn’t much slack in interpreting what the stretched elastic band tells us about Romeo. From the point of view of interpretability, there is no great difference between someone’s saying (27) and someone’s saying: (29) Romeo is so emotionally overwrought that he might suddenly do something irrational. This is not because the predicate ‘is an elastic band stretched to its limit’ comes to have a new meaning courtesy of a sophisticated theory of meaning; nor is it because there are properties in common between an elastic band in that state and Romeo’s state of mind. It is simply because the state of affairs of an elastic band stretched in the way—the very object got by semantic descent from the predicate expression taken in its ‘narrow’ truth-conditional sense—conveys information about Romeo no less efficiently than the linguistic predicate in (29).

The crucial point is that there can be uses of objects as qualifiers which do not require much input from cultural and linguistic context, and which therefore can match purely linguistic predications in respect of determinacy. However, in making this point, I asked you to consider (27) and (29) alongside one another, and I do realize that this is dangerous. It might be taken to imply that these sentences are in some sense equivalent, or even worse, that (29) is a paraphrase of (27). Neither of these implications is intended. With respect to paraphrase: I wouldn’t resist treating (29) as a comment someone might make about an assertion of (27); it is of the right form to be what I have called a rationalizing comment—a comment used

¹¹⁷ Roughly, Romeo is a loose cannon.

in justifying (27). But this is no more paraphrase than the following elucidatory comment:⁽³⁰⁾ An elastic band stretched to its limit can break at any moment, often with painful consequences. The question of whether (29) is equivalent to (27), and could therefore be a substitute for it, is separate from, and more difficult than, the question of paraphrase. It is easy to imagine someone thinking that the effect of (27), minus only its flourish, could be achieved by (29), and this raises all sorts of questions about the importance, richness, and independence of metaphor in comparison to the literal. However, I certainly didn't intend the introduction of (29) to raise these questions, and I will postpone them to the last section of the next chapter. This should not be taken as a sign of reluctance in dealing with them; it is simply that sensible answers to these questions require more detail about metaphor, and about the semantic descent account, than is on the plate just now.

4

Embellishment

Up to this point, the semantic descent account has been presented in a minimalist way. It has been motivated and defended by reference to three truths about metaphor which are I think fundamental, though they are scarcely exhaustive, and only a restricted range of examples has been used in its exposition. The aim of this chapter is an abandonment of this minimalist stance. By considering a range of further and more challenging examples, I shall show how the semantic descent account can be extended, and I shall consider a number of features of metaphor—as it were, further truths about the phenomenon—which can be naturally accommodated within my account. Given the length and detail of this chapter, and the spartan presentation of my account in the previous one, perhaps ‘embellishment’ is misleading, but I cannot think of a better title.

4.1. Complexity

Philosophical writers on metaphor often choose starkly simple examples. Though there are exceptions, the literature is peppered with sentences such as: ‘x is a pig, is a rock, is a block of ice, is a fox, is a lion, is a vulture.’ To be sure, philosophers do show some sensitivity to the fact that, in their desire to get on with the job, they appeal to what might seem to be, and often are, seriously misleading examples. But merely paying lip-service to the need for complications is not enough. Indeed, a recent writer complains, with I think some justice: One attempts to develop a theory of metaphor that accounts for these [simple] examples, which is then to be extended to cover the more complex cases. Sometimes ... such an extension will at least be sketched out. More frequently, as with Black and Beardsley, the discussion will simply break off, simply leaving it as an unexamined assumption that, somehow or other, the extension must be possible. (White 1996: 56–7)

As his notes to this passage argue in detail, White does not believe that accounts such as Black's or Beardsley's *can* be extended, or that we could ever account for the complexity of metaphor by beginning with simple examples. And I agree with him. (He himself offers an account which seems to me to suffer from the defect of beginning with examples that are too complicated. Among others, his view will be discussed in Chapter 5.) But more needs to be said about different ways in which philosophers' examples of metaphor tend to be simple, if not simplistic.

On the one hand, there is a tendency to focus on relatively simple syntactic forms; subject-predicate sentences in which the predicate counts as metaphorical are pretty well the standard. On the other hand, though often combined with this syntactic

simplicity, there is a tendency to use examples lacking the creativity, vividness, and insight shown by more realistically complex instances of metaphor. The complaint here is that philosophical discussions often turn on what are in fact examples of dead metaphor.

What the above suggests, unsurprisingly, is that complexity in metaphor is itself complicated, so I think it best to deal with its different aspects one at a time. In the next section, the topic will be syntactic complexity, and the ways in which the semantic descent account naturally extends to deal with it. Then after a brief discussion of the apparently minor phenomenon of mixed metaphor (which turns out not to be so minor after all) I shall begin to confront the kind of complexity appropriate to richer and more vivid examples of metaphor.

4.2. Syntactic Complexity

There is no obvious limit to the syntactic forms within which metaphor can be found. The following are only a sample—some of which were mentioned in passing in connection with the transparency point in Chapter 1—but they will serve initially to remind us of this variety:

- (a) Out of the crooked timber from which men are made, nothing straight can ever be built.
- (b) In cities you build a language of circumspection and tact, a thousand little intimations, the nuance that has the shimmer of rubbed bronze.
- (c) Swerving at the last moment to avoid innocent bystanders, his argument came to halt.
- (d) Her prose shows traces of the rough timbers that a more careful builder would have covered over.
- (e) He has the personality of a traffic cone.
- (f) Juliet shined when she came into the room.
- (g) ... Faunia, whose sculpted Yankee features made me think of a narrow room with windows in it but no door.
- (h) The ball I threw while playing in the park has not yet reached the ground.
- (j) When questioned, he offered his usual soap-bubble reason for what he had done.¹¹⁸

As is obvious enough, a number of these examples are more than merely syntactically complex, but I really do intend to confine discussion in this section to that most obvious kind of complexity. Certainly, there is little in the above list of the subject-predicate style which is the staple diet of philosophical discussions and, up to now, has been the only style I have myself considered. Instead, one finds metaphorically active components in verb phrases, in prepositional phrases, in the second place of

¹¹⁸ Sentence (a) is of course from Kant 1784/1912: 23 (from 'Idea for a Universal History with a Cosmopolitan Outlook'); (b) is from Don DeLillo's *Underworld* (1999: 446); (e) is simply typical journalese and comes from an *Observer Magazine* column; (g) is from Philip Roth's *The Human Stain* (2001: 207); (h) is from Dylan Thomas *Collected Poems* (2000: 152). The others were fabricated for the occasion.

dyadic predicates, in adverbial phrases, and even in a whole sentence, no ingredient of which is obviously metaphorical.

Here let me insert an aside about the expression ‘metaphorically active component’. Up to now, I have been deliberately and at one point explicitly non-committal about the use of labels such as ‘metaphor’ and ‘metaphorical’. In particular, while I have described the utterance of certain sentences as metaphors, I have not explicitly considered such questions as whether sentences become metaphors because some expression in them is used metaphorically, or whether it is more the other way around. Nor were these sorts of question pressing, given the very restricted range of examples appealed to in the previous chapters. Still, it is obvious that my account requires us to focus on words or phrases, finding in them, via descent and qualification, the source of metaphoricality. So, it might be thought that on my view a metaphor is a word or other larger, but still subsentential, element. But this is not so, as will be shown the fact that (h)—a sentence in which no subsentential item is ‘metaphorical’—is as easily handled by the semantic descent account as simple subject-predicate metaphors. The fact is that there is nothing intrinsic to my account which requires us to see it as word-based, and this seems to me right; the issue of, so to speak, the unit of metaphor has never been a fruitful one. So, I will continue to be non-committal on the application of the terms ‘metaphor’ and ‘metaphorical’, even though I obviously have to employ them, sometimes of words and sometimes of sentences.

Any account of metaphor must face up to syntactic diversity, it must explain how it is possible that metaphorical elements can figure in virtually any grammatical slot. This might seem a large enough task, and it is certainly a task too weighty for most of the accounts canvassed in Chapter 1. But, while squaring up to this diversity, it is important not to lose sight of an underlying feature of metaphor, one which I believe to be responsible for philosophers' otherwise irresponsible obsession with subject-predicate examples.

Perhaps because of, but certainly since, Aristotle's seminal discussion, we have come to think of metaphor as suggesting some kind of transference from one subject matter to another, typically a comparison, often unexpected, between one sort of thing and another. That there are such comparisons cannot be denied, but I have already denied that comparison itself can be the basis for an account of metaphor (and I will reinforce this denial later in this chapter and in the next one). Still, even if comparison falls short of what is required in a full account, it does reveal a kind of bipolarity that is genuinely typical of metaphor. Moreover, it is this bipolarity that probably lies behind the ease with which, in the face of so much contrary evidence, writers treat the ‘S is P’ form as paradigmatic. It seems therefore reasonable to require of any account *both* that it handles syntactical forms other than ‘S is P’ *and* explains, or at least explains away, our sense of the fundamentally bipolar—‘this-is-that’—nature of metaphor. Unsurprisingly, I think that the semantic descent account fulfils these requirements.

4.2.1. The shining

Each of the elements that make up the semantic descent account have crucial roles to play in dealing with the twin requirements of syntactic diversity

and the intuition of bipolarity. On the one hand, it will be necessary to enrich our understanding of the idea of 'linguistic control' as it figures in the move from word to object in semantic descent. And, on the other, it will prove vital to take seriously the diversity of objects that we employ in qualifications—a diversity I have insisted on, but which was not evident in the narrow range of examples used in Chapter 3. However, rather than go on about these two aspects in a general way, let me illustrate what I mean by starting with one of the 'easier' examples in my list: (f) Juliet shined when she came into the room. Admittedly, the metaphor here is not particularly vivid, but that is not relevant to the narrower issue of syntactical complexity under discussion. Moreover, whatever its defects, the metaphor in (f) is not obviously one we would think of as 'dead'. Assuming that, as it is used, we have identified (f) as a metaphor, the main burden of responsibility for the metaphorical effect falls naturally on the intransitive verb 'shined'. Marking that verb with the indicators of semantic descent we get: (f1) Juliet ↓shined↓ when she came into the room. And, following the previous use of this notation, the next step should be the search for some object which we can count as a qualifier—a non-linguistic object which figures in the extension of 'shined' and to which we are led by these markers of semantic descent.

Now it is not obvious that verbs like 'shined' have extensions, and I could imagine all kinds of resistance to a semantic theory which insists that they do. However, we do not need to be committed to any very heavy semantical machinery to understand the point of the down-arrows in (f1); familiar ways of thinking about verbs will serve well enough. Thus, if we are told that the tide turned, it is natural enough to think this is made true by the occurrence of some particular event, an event in which the tide figures as the agent (or subject). When asked by someone to explain the changed orientation of boats at their moorings, one can say 'The tide turned'. But one can also, and no less naturally, say that it was the turning of the tide that caused the boats to change their orientation. Moreover, it is scarcely controversial that in saying this latter thing, we are picking out, or referring to, some event. And, just as we find it natural to refer to the event of turning in this case, we can for convenience, and ignoring its rebarbateness, speak about the expression 'shined' in (f) as picking out or referring to some *shining*. In saying this, we are saying no less, and no more, than is intended by notation in the proto-predicative '↓shined↓'.¹¹⁹

To be sure, just as there are lots of different kinds of turning, dependent for their different character on the agent involved, so there are lots of events which are shinings. It is at this point that a certain enrichment in the idea of linguistic control enters the

¹¹⁹ I do not mean here to be taking issue with the Davidsonian insistence that 'The tide turned' existentially quantifies over events rather than referring to some particular one. Aside from the fact that Davidson's interest and mine are quite different, his account is compatible with what I say. He wouldn't deny that, in a certain context, we might be interested in the specific event that by his lights is not strictly referred to by some relevant sentence, but which nonetheless makes that sentence true in that context.

picture. Clearly enough, the word ‘shined’ in (f) exercises this much control over the descent: the event picked out must be one in which there is an agent or subject, but no object. After all, ‘shined’ is intransitive. Additionally, like any utterance, (f) will have been uttered in a context, whether one set up linguistically and/or in other ways. Thus, one could imagine it as occurring (with apologies to Shakespeare) at some point just after the initial speech in the balcony scene. Or as said by someone who has been present during Juliet's careful preparations for making a grand entrance. (Imagine witnessing her assiduous routine of hair brushing, facial scrubbing, the application of creams and makeup, etc.)

Depending on these different circumstances, one might well be led to think of the shining event as having either the sun or some lovingly polished surface as its agent, one or the other of these being the event on which the burden of qualification falls. As events, these are certainly distinct from material objects like the sun or infants that were thought of as qualifiers in the previous chapter, but in having some definite subject, these events share a certain determinacy with those sorts of object. This determinacy is important, but it should not be confused with the fact that in certain examples—in particular the headline example of Juliet being said to be the sun itself—the qualifying object is an existent object.

Before the phenomenon of metaphor even entered the picture, that is, in the cases of qualification described in Chapter 2, the whole point of examples was to display the ways in which objects took it upon themselves to play a role usually accomplished by words. Given this, it is obvious enough that any object I used to illustrate qualification must have been not merely determinate, but existent. In order for an object to be salient enough in some context to take on a predicative function, that object must be present (or perhaps just recently present). *But qualification is not metaphor.* Metaphor is a linguistic phenomenon, it is something we accomplish with words. To be sure, predication is also a linguistic function, and the business of Chapter 2 was to investigate the possibility that this function could be taken over by objects. However, though it is an easy mistake to make, it is not here my suggestion that the linguistic function of metaphor is also something that can be taken over by objects. The aim is instead to show how the predicative function of objects, that is, qualification, can be *part of* an account of metaphor.

In metaphor but not in qualification, we begin with words that, like ‘shined’, are perfectly ordinary. Precisely because metaphor begins with words, it allows not only what I have called ‘linguistic control’ over semantic descent, but in addition a certain kind of imaginative playfulness that is not available in primitive cases of qualification. This is not to deny a role for imagination in qualification: the juxtaposition of objects characteristic of qualification undoubtedly calls on that faculty. It is just that the presence of words in metaphors allows room for a certain kind of exercise of the imagination that is not available in primitive qualification. This is evident even in a relatively straightforward case like (f). That sentence speaks of a shining, and we all have some idea of what that kind of event is like. But, if I am right about the need for the semantic descent that this sentence encourages, we are invited by (f) to think of some quite specific or determinate shining event, one that it is plausible to

regard as a qualifier of Juliet. We are being invited to imagine that there is in fact some such event—whether it is the sun's shining, that of some carefully polished surface, or some other—which we could have witnessed, even though no such event is actually taking place. This is precisely the playful use of the imagination that I suggest is ubiquitous in metaphor, but not in qualification by itself. In metaphor, we are free to conjure up objects which, because they are imagined cannot be existent, but which can nonetheless be quite determinate.¹²⁰

An important aside: what about an event of Juliet's shining? Mightn't this be the event that (f) speaks about? Well, the simple answer here is that the words in (f) simply do not conjure up such a thing. Not only is Juliet not shining, it does not make the least sense to say that she is. To think otherwise is to be wildly misled by superficial grammatical appearances. To be sure, many have been and continue to be misled. There are hordes of philosophers who regard (f) as somehow making this claim about Juliet, albeit falsely. But, construed as they suggest it should be, (f) is more a piece of arrant nonsense than a false claim. Would we stand still for a view which insisted that the notorious 'Green ideas sleep furiously' is false? Well, (f) as understood by the false-sentence brigade is no less problematic.¹²¹

That our imaginations do actually conjure up a determinate event referred to by—or, if you are squeamish about using this term, corresponding to—'shining' is both evidence *of* and evidence *for* the semantic descent account. To see why, begin with this thought: we need there to be determinate events (objects) in order for the qualificational part of the semantic descent account to come onto the scene. By itself, this is scarcely proof of the account, sounding rather more like a blatant example of *petitio*. But suppose there to be independent reason to think that we do conjure up appropriate determinate events in our dealing with metaphors like (f). Then, since the semantic descent account has a job of work for these objects—namely as qualifiers—the semantic descent account promises to explain this otherwise surprising employment of our imagination. After all, why would we bother to imagine a determinate object in respect of (f) unless there was a real purpose that it is intended to fulfil? The sentence itself appears to speak of Juliet's shining—something that is about as intelligible as the repose of a furious green idea—and the actual words in the sentence make no mention of anything else which would lend determinacy to the event. So, if we do in spite of all this find reason to think that (f) conjures up a determinate object—an event of shining—then this does call for some explanation.

¹²⁰ I realize in this speaking of exercises of the imagination, and exercises of this or that faculty, it sounds as though I am engaged in precisely the psychological enterprise that I earlier eschewed. And this leaves many of the things I say in this chapter open to worries about whether we are aware of what we are doing in processing metaphors or aware of the intentions with which they are offered. However, all of my comments could be re-presented from the perspective of the theorist attempting to account for the intelligibility of this or that metaphor or phenomenon of metaphor. I haven't consistently adopted this perspective simply because it enormously complicates the exposition.

¹²¹ Perhaps alone amongst recent writers, White is very good on this subject (see especially Ch. 2 of White 1996). Note too that for all that the false-sentence story is just plain wrong, it can seem so plausible simply because, if I am right, we do conjure up a shining event and we do use (f) to make a claim using this event.

Evidence that we do actually think the shining event only touched on in the words of (f) is actually quite determinate is shown graphically by the kinds of thing we find it natural to say (or think) ‘downstream’ of the utterance. Thus, if uttered in what we can call the ‘Romeo-context’, one could easily imagine (f) being followed in speech (or thought) by: (f2) Juliet positively dazzled the guests. (f3) When Juliet saw her rival, a cloud seemed to pass over her. (f4) By the end of the evening, she was well on the way to setting. Or more waspishly still: (f5) She had us reaching for our factor 15 to avoid skin damage. And, in another context, continuations might include: (f6) One could see your reflection in her face. (f7) There was no hint of the hard work that had made it so. It is clear enough that these are all perfectly intelligible as comments on (f), and their very intelligibility leans on the recognition of some determinate kind of shining already at work in its utterance. Let me spell this out.

When we hear (f), I suggest that we let our imagination fill in some particular shining event, a choice constrained by the context, including here the linguistic context, in which (f) is uttered. In this way, when we do hear the above kinds of continuation, we are well prepared for them. An aside: It might be held that the determinacy only comes into play when such continuations are actually offered—that they function to *extend* the metaphor in (f) precisely by making further comment about it. Later in this chapter, I will have more to say about the idea of extending a metaphor but for now I note simply how implausible it is to regard the continuations in this way. If we had not been thinking that, for example, it was the sun's shining at work in (f), then what exactly is it we are ‘extending’ when we say things like (f2)–(f5)? The fact is that there is nothing there to be ‘extended’ unless, as I suggest, we take the metaphor in (f) to call on the determinate object—the sun's shining in its characteristic way—which qualifies Juliet's entry into the room.

Clearly enough, the event conjured up by (f) is not one actually taking place, and is therefore not one that can be demonstrated. It thus does not serve as a qualifier in precisely the way that, for example, Nabokov's *paling* did. But this is no more than a confirmation of what I have already insisted on, namely that metaphor is not itself an exercise of what I have been calling ‘primitive’ qualification. Having a linguistic starting point, metaphor allows a degree of play to the imagination that is simply not possible otherwise. Still, I do insist that the objects conjured up by the imagination are there to fulfil precisely the same function as those more palpable items used

in primitive qualification: they serve to provide information, in the case of (f) about Juliet's entry in a room, and are as determinate as is necessary for this way of understanding (f) and its various continuations.¹²²

Just before turning to the other examples in my original list, let me insert a note about bipolarity. The sentence (f) is not of the subject-predicate form, and the treatment suggested for it by the semantic descent account does not depend on its being twistable into that form. Yet there is a sense in which the 'this-is-that' cadence of metaphor can nonetheless be heard in this treatment. Thus, we begin by recognizing that the verb in (f) allows descent to an object, in this case, an event or action. As noted above, the event-object might be that of the sun's shining, or some surface's shining, both are certainly possible, and the choice will turn on linguistic or non-linguistic contextual factors. Moreover, while the determinacy of being an event of a certain type, with a certain agent, is important, this event-object need not actually exist anywhere but in the imagination of the interpreter of (f). All that said, when one does confront an event, say, that of the sun's shining out of a cloudless sky, it is this event-object which qualifies, not Juliet, but Juliet's entry in the room. It is the imagined (but determinate) event of the sun's shining which serves as the source of information about the event which is Juliet's actual entry. In short, we have something of the 'this-is-that' form, though that we have this does not depend on the syntactically banal examples so often appealed to in philosophical discussions.

As it happens, the obvious way in which to redraft (f) against the background of the semantic descent account would yield something of the subject-predicate form: Juliet's entry into the room was the sun's shining in a cloudless sky. And perhaps surprisingly this is an intelligible and rather genteel way of putting what is in effect the same sentiment as the original. But that this transformation works in this case is no guarantee that it will always do so. The underlying 'this-is-that' nature of metaphor is ubiquitous, but it is simply not the case that we can always find a subject-predicate formulation which works as well as that given above. Moreover, there are interesting and important reasons for this which will become apparent as we move on to some of the other examples in the original list. Still, the very fact that we can construct such a sentence, based as it is on the semantic descent account, suggests that the account can accommodate syntactic

¹²² White's account of metaphor (1996: Ch. 4) includes much that is relevant to—though still quite different from—my discussion of imagination and determinacy. He too finds that there is an underlying imaginative addition to the surface forms of sentences like (f), but he sees it, not so much as a matter of pretending there to be a determinate shining event, but rather as the provision of a 'secondary' sentence which is straightforwardly literal, and whose vocabulary gets mixed in with a version of the original that he calls 'primary'. Thus, he would see (f) as somehow an amalgam of two sentences: one closer to the original (f) and one that might read: 'The sun shined in a cloudless sky.' I discuss White's very interesting view in the next chapter, though various of his specific observations about metaphor will further figure in this one.

forms other than subject-predicate, as well as the bipolarity so often associated with metaphor.¹²³

4.2.2. Decisions

The next sentence to be considered:(d) Her prose shows traces of the rough timbers usually concealed by more careful builders, is syntactically more complicated than (f), but that fact alone is not my primary concern. Instead, I should like to focus on two special features of (d): first, that it requires us to recognize a hitherto unnecessary complication in the semantic descent account; and, second, that it is one of those metaphors which, for all that it shows the usual bipolarity of metaphor, is not really convertible into subject-predicate form.

Let me begin with a fairly obvious suggestion for annotating (d) with down-arrows:(d1) Her prose shows ↓traces of the rough timbers usually concealed by more careful builders↓. The phrase ‘traces of the rough timbers ...’ is marked for semantic descent, and we are thereby required to think of something in the extension of this noun phrase. This is straightforward, though note that in doing so one is bound to be led to think of the traces as featuring in some wall, ceiling, or similar structure, since traces of rough timbers are most naturally thought of as features in some such structure. Hence, what we are led to by the noun phrase marked in (d1) is a determinate state of affairs, a way it is in regard to some wall or ceiling that has been worked on by a careless builder. And it is how it is with respect to this wall or ceiling—some determinate state of affairs—that we are invited to think of as qualifying someone's prose style.

One could dramatize the circumstances which led to (d). You are writing a review of a book. Distracted from your task, you notice your recently constructed study wall, replete with the annoying bulges where the underlying 2 3 4 timbers show through. Just like the prose style of this author, you think, and you return to work by typing (d), a sentence you intend to slip into the review in an appropriate place.

In whatever way it comes to be uttered or written, there is an important feature of (d) that is not shared by (f). Whereas there is only one obvious way in which to

¹²³ There is a nice discussion in White 1996 about the relationship between Black's focus and frame terminology and Richard's tenor and vehicle. The latter is given full marks; the former is treated as hopelessly inadequate. This can seem surprising, especially to those who regard the two pairs as simply alternative vocabulary. White bases his differential attitude on the fact that Richard's is offering a comment about metaphorical *thought*, whereas Black's is about the *linguistic form* of metaphor. But while this distinction is not by itself all that clear or convincing, I suggest that the comments in the text about the underlying bipolarity evident in the semantic descent account—especially when put alongside the resistance of many metaphors to being twisted into subject-predicate form—vindicate White (and Richards). I should also note here that ‘bipolarity’ is White's term.

mark (f) with down-arrows, this is not so of (d). The representation in (d1) is plausible, but so are (at least):(d2) Her prose ↓shows traces of the rough timbers usually concealed by more careful builders↓,(d3) Her prose shows traces of ↓the rough timbers usually concealed by more careful builders↓.While not differing radically, these alternatives are different: they do not count all the same words in the original as relevant to the metaphorical effect of the whole. In terms of the semantic descent account, this differential parsing of words leads to real, if sometimes subtle, differences in our understanding. If we go for (d2), we are encouraged to think about the author's prose in terms of something that is accomplished, as it were, by the wall (some showing of features), whereas in (d3) we are encouraged to use a material object (timbers in a wall) in characterizing what the prose itself does.

Subtle though they are, the fact of these differences demonstrates something that is true of (d) independently of the semantic descent account, namely that the vocabularies of metaphorical sentences cannot always be divided into those which are intended literally and those which active in the metaphor. For obvious enough reasons, no such problem of division showed up in 'Juliet is the sun' or in 'Juliet shined when she entered the room', but once you think about sentences like (d) you can see that this phenomenon is ubiquitous, and though it has tended to be overlooked, it is actually a rather important feature of metaphor. White (1996) called it 'bifurcation' and found its general neglect to be one of the central charges against many philosophical accounts of metaphor. In this I think he was certainly right. But it is one thing to isolate some phenomenon of metaphor and another to explain it, and here White's own account seems to me deficient. (I will discuss White's own proposal in Chapter 5, but what I want to stress now is the natural way in which the semantic account deals with what White (not wholly perspicuously) called 'bifurcation'.)

Though it will vary in different examples and in different contexts, the explanation of this division in vocabulary offered by the semantic descent account is simple: if certain words of the metaphorical sentence are drafted into the process of descent, then they figure as part of the metaphor. But if these same words are left out of this stage of the account, they can still play a role in guiding the qualification effected by the object to which descent is made. Thus, as marked in (d1) what the prose *shows* is something we come to understand by focusing on the *traces* of the rough timbers; in (d2) it is the *prose itself* that is qualified, here by *something shown in some structure*; and in (d3) it is what the prose *shows traces of* that is qualified by the *structure itself*. What one can see in these different formulations is a trade-off between the vocabulary of semantic descent and the vocabulary of qualification. Even more importantly, even the simple case of (d) shows how there can be an

oscillation between these different versions—an oscillation which makes possible effects not achievable otherwise.¹²⁴

Go back to the scene in which I imagined you to be writing a review: the uneven wall was on your mind, so was the annoying prose style of the author. Suppose that you had tried to connect them by saying or thinking:(d4) The author's prose is a wall in which one can see the bulges made by the timbers ... This flat-footed formulation is not merely weak, it is just not suited to its purpose. The basic problem is obvious enough: directly juxtaposing prose and a wall brings us up short; it is far from obvious that there is an intelligible connection between words and walls, except perhaps the irrelevancy in the cliché 'wall of words'. Thus, though something about the wall seems to convey something about the prose, putting the one in direct contact with the other just doesn't work.

One could try this:(d5) The author's prose is like a wall ... This sounds better than (d4). The use of the word 'like' usually marks comparison or, in the right circumstances, simile, but it can also have the effect of softening a rather too far-fetched direct comparison. For now though I shall put 'like' on one side: a discussion later in the chapter will take up the vexed issue of simile and metaphor, and at that point I will suggest something rather radical both about that issue and about the role of 'like' itself. In any case, we already know a better way to formulate the wall–prose metaphor that doesn't need to call on the rather blunderbuss device of explicit simile.

What needs to be done is to somehow hide the wall, so to speak, so that it does not come too directly into contact with the prose, and of course this is just what is achieved in (d). As the various semantic descent alternatives show, (d) keeps the wall (or, in any case, whatever is the imagined and determinate structure) implicit: the words which oscillate between metaphorical and literal treatments create just the right separation. Thus, it is not the wall that is used to qualify the prose, but the *way* it is with the wall and what the prose *shows*. Or what the wall *shows* and the way it is with the prose. Or the *traces* of something in the wall that capture features of the prose. Or the wall that captures something about *traces* in the prose. No one of these is right to the exclusion of the others, but the fact that there is this degree of play in the semantic descent account of (d) is just what is needed to keep prose and wall apart, while still allowing them to work together in the metaphor.

¹²⁴ In the previous chapter, I noted a kind of vibration or resonance between the linguistic form 'the sun' and the sun itself. This resonance helped to explain transparency, while at the same time keeping faith with the idea that the literal is alive in the metaphorical. At a minimum, the oscillation described in the text is merely another form of this kind of resonance: we move back and forth between certain linguistic forms and certain objects. However, it has an added feature: one not only moves from the linguistic to the non-linguistic, but also from one non-linguistic object to another. It is thus a more complex kind of resonance—one whose special importance will be further emphasized in the text—and that is why I reserved the term 'oscillation' for it.

Note too how difficult it is, not to say impossible, to twist (d) into a subject-predicate form that preserves any of the original metaphor. Attempts to do this, as in (d4), are bound to bring the wall into too direct contact with the prose. Still, when (d) is given the treatment I recommend, then no matter how we parse the words in it for descent, we can clearly discern bipolarity. This is because, in achieving semantic descent, we are forced to find some determinate object—whether a material one like a wall or a state of affairs such as what a wall shows traces of—which is the source of information about the prose or prose style. My guess is that any metaphor which allows the kind of oscillation found in (d) will not lend itself to re-statement in subject-predicate form. Some further evidence for this will be given shortly, but nothing crucial hangs on my being right about this. For the real contribution of semantic descent is that it allows us to deal with forms that are definitely not subject-predicate, and which may well not be twistable into that form. Moreover, even while coping with a variety of syntactic forms, it shows why there is nonetheless a universal or near-universal underlying bipolarity in metaphor.

That the phrase ‘shows traces of’ can be either wholly or partly absorbed into either the literal or metaphorical part of (d) might suggest that we should regard it as apart from the literal ‘her prose’ or the metaphorical ‘the rough timbers ...’. Perhaps we could even count it as doing what ‘like’ can do, but in a more specific-to-context way. Neither of these options are unreasonable, but they are not forced on us. The semantic descent account shows how the phrase can be absorbed intelligibly, either partly or wholly, into either functional unit, and the oscillation that is set off by the different ways of achieving this absorption is perhaps enough to suggest that these words are special. Moreover, there is a reason for resisting the treatment of phrases like ‘shows traces of’ as essentially outside the metaphor/literal divide. Quite simply, there are cases in which bifurcation cannot be handled this way.

An example that figured in Chapter 3 makes the point nicely. Though Davidson's version was the simpler, what the critic actually said of Tolstoy was: (t) Tolstoy is a great moralizing infant. Focus here on ‘great’: does this function literally or metaphorically in this sentence? Semantic descent can offer these two versions¹²⁵: (t1) Tolstoy is a ↓great moralizing infant↓. (t2) Tolstoy is a great ↓moralizing infant↓. As with (f), there is thus a choice: we can consider ‘great’ as a modifier of the qualifying object—the infant—or of Tolstoy himself, and there can be an oscillation between these two readings. We are certainly not called upon to decide in favour of one or the other, and both are accommodated by the semantic descent notation. However, unlike the case of (f), there is no special problem about subject-predicate

¹²⁵ Obviously enough, there are issues of bifurcation in regard to ‘moralizing’ as well as ‘great’, but it will simplify present discussion to deal with only one of these bifurcated words. ‘Moralizing’ makes the sentence more complex and interesting, but doesn't add anything to the basic point made in the text.

form: as it is already in that form, we do not need help to understand how (t) can display the ‘this-is-that’ character typical of metaphor.

Given that the oscillation between the two readings of ‘great’ is not needed to finesse an otherwise awkward juxtaposition, its point must lie elsewhere. Nor is it difficult to see what it is. For it should be clear enough that the variant readings of ‘great’ add a certain richness, even to this most simple example. The oscillation of ‘great’ in (t) conveys at a stroke both the genuine human and literary *greatness* of Tolstoy and his nonetheless *extreme* childlikeness. That there are these two readings intensifies the point made by each of them, rather as if there was a kind of ‘micro-metaphor’ revolving around this single adjective.

Bifurcation no doubt serves specific purposes other than the two discussed above, but they seem to me to be the main types. Trade-off between literal and metaphorical reading, and the oscillation thereby engendered, functions either to make palatable what would otherwise be an awkward comparison, or to enrich a comparison that might be too bland. (Think here of the difference between (t) and Davidson's ‘Tolstoy is an infant’, even without ‘moralizing’.) In considering some of the other examples on my original list, I shall have reason to return to the first of these functions, but just before that, something must be said about a notational novelty needed to cope with:(c) Swerving at the last moment to avoid innocent bystanders, his argument came to halt. So far, my notation has only been called upon to highlight words between left- and right-hand down-arrows. But the above sentence cannot be dealt with so straightforwardly. The words in (c) which need to be marked for semantic descent seem to bridge others, which are arguably non-metaphorical. Thus, I can imagine the following notation:(c1) ↓Swerving at the last moment to avoid innocent bystanders↓, his argument came to a halt, but I can equally well understand that one might want to include ‘came to a halt’ as figuring in the process of semantic descent. After all, it is the swerving *and sudden stopping* which serves here to qualify the course of someone's argument. Easily accommodated in various ways, here is my suggestion:(c2) ↓↓Swerving at the last moment to avoid innocent bystanders↓, his argument ↓came to a halt. ↓↓Adapting the familiar use of quotes within quotes, the first *pair* of down-arrows marks the beginning of the semantic descent phrase, there is a temporary pause at the single down-arrow which is taken up again with the second single down-arrow and the whole terminates with the final pair. The phrase thus marked, namely ‘Swerving at the last moment to avoid innocent bystanders ... came to a halt’, conjures up an event involving a vehicle scarcely under control, but just managing to

stop without doing any damage.¹²⁶ (If the context is philosophical, then perhaps a madly ridden bicycle would do the trick.)

Note that versions (c1) and (c2) are both reasonable: as with previous examples, there is the possibility of trade-off between semantic descent and qualification, a trade-off which generates an oscillation between two readings. In (c1), the argument's coming to a halt—that event—is qualified by the mad swerving of a vehicle only just under control; whereas in (c2) the argument itself is qualified both by the way in which in the vehicle was driven and the manner of its finally stopping.

Can the sentiment in (c) be captured in subject-predicate form? A moment's thought should convince you that it cannot be. As in (d), there is an awkwardness in juxtaposing the two obvious items that figure in the metaphor. We can scarcely capture the idea of (c) by saying: (c3) His argument is a bicycle that swerves at the last moment ... So, the apparent oscillation between versions—one in which some argument is the subject, and one in which the subject is the way that argument came to a halt—acts as a buffer between the two objects. That the words 'came to a halt' are variably interpretable serves the same function as 'like' does in certain contexts. Using 'like', we could even fix up (c3) above: it is surely somewhat more palatable to say that the argument is *like* a bicycle swerving madly than that it *is* such a bicycle.

A final comment on (c): unlike the prose-wall example (d), and more like the Tolstoy example (t), the oscillation evident in (c) serves to intensify the metaphorical effect. The very fact that the halting is now that of the metaphor 'vehicle' (punning on Richard's nomenclature, perhaps unforgivably), and now that of the argument itself, sharpens our focus on what is actually being said about the argument. Thus, though I have only made the point briefly, in (c) we have a case in which oscillation fulfils both of its most obvious functions simultaneously.

4.2.3. Like '*like*'

I have taken a lot of trouble over only three of the metaphors on the list given at the beginning of this chapter, but things can move more quickly now. For the ways of coping with the syntactical complexities of (f), (d), and (c) will allow us to sweep up three more of the examples at one go. These are:

- (b) In cities you build a language of circumspection and tact, a thousand little intimations, the nuance that has the shimmer of rubbed bronze.
- (e) He has the personality of a traffic cone.
- (g) ...Faunia, whose sculpted Yankee features made me think of a narrow room with windows in it but no door.

¹²⁶ White (1996) would have us put a kind of variable subject in place of the ellipsis in this phrase, e.g. an X. This is because he insists on finding in every metaphor a pair of complete sentences. I don't follow him in this, since what matters for me is simply that there be a form of words which conjures up an object that can then serve as a qualifier. In (c), especially on the reading (c2), the object is an event, so a sentence-sized form of words can seem appropriate. But this will not always be the case, and when it isn't White's insistence on sentences is a liability.

Sentences (b) and (e) share a common feature: they both use the phrase ‘has the F of’, where what replaces the ‘F’ is something that can feature either as part of the description of the qualifying object or as part of the description of the subject of the metaphor. This is easiest to see in the simpler (e), a sentence which is admittedly intended as no more than a joke, but is no less a metaphor for that. We can mark that sentence for descent in either of the following ways: (e1) He has the personality of ↓a traffic cone↓. (e2) He has ↓the personality of a traffic cone↓. where it is either the traffic cone that is said to qualify some individual's personality, or it is the personality of the traffic cone which qualifies the individual himself. As with the earlier examples, these multiple interpretations, far from being hindrances, are actually integral to understanding the sentence correctly. None of the straightforward subject-predicate versions work: (e) is not captured by ‘He is a traffic cone’ or ‘His personality is a traffic cone’. There is a way to twist (e) more or less into subject-predicate form, namely: (e4) His personality is the personality of a traffic cone. But of course this displays in a graphic way the dual role of ‘personality’.

As already suggested, the key here is ‘has the F of’. This form of words is common in metaphors whose syntactic structure is more complicated than ‘S is P’, and it always involves the kind of bifurcation just noted. The semantic descent account deals with this by the kind of trade-off and oscillation described earlier. One can see ‘personality’ either as what is qualified or as part of what is needed to do the qualifying, and the dual role evident in (e1) and (e2), as well as in (e4), smoothes over what would otherwise be an awkward juxtaposition.

In effect, ‘has the F of’ serves the hedging function we saw earlier with ‘like’, though it is more specific to context and informative than a use of an unqualified ‘like’. So, while ‘He is a traffic cone’ is simply bizarre and ‘He is *like* a traffic cone’ is not bizarre but is unclear, the sentence: He is like-in-personality a traffic cone, is just about right; it is more or less the same as (e).

Example (b) is superficially more complicated than (e), but it involves the same ‘has-the-F’ form of words. There are two main differences: first, there is a preamble which spells out the nuanced something that is the literal subject of the metaphor; second, the word that plays the role of F, namely ‘shimmer’, goes more naturally with the qualifying object (rubbed bronze) than it does with the literal subject, the social life of cities, that is, ‘the language of circumspection and tact, a thousand little intimations’. Let me expand on this second difference a little.

While we fully understand *literally* the idea that an individual human being has a personality, we are pushed when it comes to attributing personality to a traffic cone. But of course being pushed in this way is crucial to the metaphor. ‘Personality’, singled out as it is by the very fact that it has a dual role, is pivotal in our being able

to make sense of a traffic cone's being used as a qualifier of an individual human being. In (b), 'shimmer' also plays a dual role, but we need to be pushed here, not to see it as applied to the qualifying object—the surface of a rubbed bronze—but as applied to the literal subject of the metaphor. And, whether the push is towards or away from the literal subject, the role of the item filling the 'F' place in 'has the F of' is crucial to smoothing over what would otherwise be an unintelligibly awkward juxtaposition.

Example (g) contains another common way of getting around the awkwardness of immediate, subject-predicate, juxtaposition, one that does not involve bifurcation, but is rather more explicit. Perhaps we can just about make sense of: Faunia's sculpted Yankee features are a narrow room with windows but no door. But when the author makes explicit reference to a process that could have led to the metaphor—one thing 'made me think' of this other—the awkwardness disappears. Moreover, this explicit reference is pretty direct evidence for the semantic descent account. That said, however, a degree of caution is needed.

Clearly enough, lots of things can make us think of other things. I see my watch lying on the bedside table, and I remember that I am to make an important telephone call just after breakfast: the watch made me think of the call. However, this kind of cross-referential use of 'things' doesn't get to the bottom of what is going on in (g). In that sentence, it is not merely that one thing refers to another, it is that one thing characterizes the other. When the narrator says that Faunia's features made him think of a certain sort of room he is telling us that information about Faunia and her features can be got from the room. The 'made me think' here is not merely cross-referential, it is can be taken as an explicit acknowledgement of the process of qualification that is fundamental to the semantic descent account.¹²⁷

If I am right about this, then it should be possible to use the expression 'made one think of' (or near equivalents) in every metaphor. And a moment's reflection confirms that this is possible. Thus, instead of 'Juliet is the sun', we can have 'Juliet made me think of the sun'. Or, in a more complex case, we could have 'Her prose made me think of a wall in which one could see timbers that a more careful builder would have covered over' in place of the original. To be sure, there can be loss in this degree of explicitness. Shakespeare's revised sentence limps after this treatment, just as it would if he had written 'Juliet is like the sun'. But any such loss comes from what is sometimes an intrusive and unnecessary explicitness, rather than from a substantial departure from the simpler metaphor.

¹²⁷ A further note about 'made me think': one could be, but should not be, seduced into thinking that this phrase marks a causal relationship between the room and the information about Faunia it provides. Roth invites us to treat the room as characterizing Faunia, and this is a wholly communicative intentional act on his part. But this effect is intensified by the suggestion given in the phrase that Roth's narrator somehow had already experienced the room as imposing this information on him, rather as the paling imposed itself on Nabokov.

There is another variation on ‘made me think of’ which seems to me to have become widespread in the relatively recent past: this is the expression ‘is a metaphor for’, as in this sample plucked from a piece about a recent shipwreck in the Bay of Biscay: The sunken tanker and its unpredictable cargo which might devastate the coast at any time is a metaphor for the terrorist menace facing Western nations. As with ‘made me think’, the expression ‘is a metaphor for’ *can* be an explicit marker of what is by my lights the use of one object as a qualifier of another.¹²⁸ Like ‘like’, these expressions often smooth over, or even make possible, what would otherwise be an awkward, or even bizarre, juxtaposition in the simpler metaphor form ‘S is P’.

4.2.4. Finishing the list

The two sentences at the end of my original list represent extreme ends of the scale of syntactical complexity. In: (j) When questioned, he offered a soap-bubble reason for what he had done, a single adjective carries the weight of metaphor, and in: (h) The ball I threw while playing in the park has not yet reached the ground, the words are all quite literal, though when you realize that this sentence has a subject that is offstage, the inevitable conclusion is that the whole sentence is somehow a metaphor.

In ‘Juliet shined when she entered the room’, we have a sentence in which a single word carries all the metaphorical weight, but, as that word is the main verb, there is a tendency to describe the whole sentence as a metaphor or as metaphorical. This is not how it is with (j). Here the sentence *contains* a metaphor, though the whole sentence wouldn’t ordinarily be thought to *be* one. One way to think about (j) is as containing, in the adjective-noun combination ‘soap-bubble reason’, a compressed subject-predicate metaphor: His reason is a soap-bubble. But this is a clumsy way to deal with this syntactic form, and there is no need for it. Using the notation of semantic descent, we can display (j) as: (j1) When questioned, he offered his usual ↓soap-bubble↓ reason for what he had done, and we can then treat the soap-bubble—that very object—as the adjectival modifier (more accurately: qualifier) of the reason. Indeed, this suggestion is a reprise of the much earlier discussion of exemplification and predication in Chapter 2. When someone displays a colour sample, and says: My house is this blue,

¹²⁸ I emphasize the ‘can’ because, as with the expression ‘made me think of’, the phrase ‘is a metaphor for’ sometimes indicates what is merely a kind of cross-reference. More on this later in the chapter.

the demonstrative certainly refers to the colour sample, but that is not all there is to it. For the colour sample thus demonstrated fulfils the same function as, for example, the adjective 'light' in 'My house is light blue'. Whereas in the latter case we have a purely linguistic adjective-noun combination, in the former we have an object functioning adjectivally in combination with a noun. And this is precisely what we have in (j1), though here, unlike the colour-sample case, the adjective-noun phrase that led to the qualification is metaphorical.

In (h), we have a whole sentence that makes a perfectly straightforward (even if unlikely) claim. No word or phrase in it is, as one says, metaphorical, yet there is little doubt that we have a metaphor on our hands. (More about this below.) These features make (h) particularly troublesome for most accounts of metaphor, but not for semantic descent. We begin by extending the notation of descent to the whole sentence: (h1) ↓The ball I threw while playing in the park has not yet reached the ground↓. As per the semantic descent indicated by the down-arrows, what should come before our minds, so to speak, is a determinate event (though admittedly one it is not easy to think of as actual): a ball thrown by a 5-year-old child, arching perhaps but never falling as years go by. Now, as that event is one described by the sentence (h), one might wonder whether there is any real need for the down-arrows; the author is telling us about an event by using (h), so is it really necessary to rub our noses in it by using this notation? And certainly one might wonder why this odd but otherwise intelligible literal sentence is given this treatment. However, these very questions serve to underline (not undermine) the semantic descent account. For while (h) is like any other sentence that describes an event, one must remember that it is not only the words making up the sentence which figure in the metaphor; it is the object brought to our attention by the those words. For that object—in this case an event—comes to serve as a qualifier of the metaphor's subject. And what is that subject? Come to that, what tells us that this is in fact a metaphor in the first place, and as such a candidate for semantic descent?

The second of these questions is about the identification of metaphor, one I have explicitly avoided, not because I have no answer, but because I don't think that any single answer is likely to be true or informative. Still, (h) nicely illustrates the complex and theoretically mixed way in which metaphors come to be identified. First off, the implausibility of the event described is bound to alert us to something out of the ordinary. But second, and implausibility aside, we might be influenced by the very fact that the kind of thing described is very ordinary, even though it occurs prominently in a poem. Poetry is certainly a locus of figurative language, and the tension between the poetic context and the apparent banality of the observation suggests metaphor.¹²⁹ There are other features, those arising from the surrounding linguistic context, which

¹²⁹ There is some evidence that we look for metaphors merely because we believe the source of the text to have been a poet. Gibbs, Kushner, and Mills (1991) suggest that the certain phrases are treated as metaphors when subjects believe them written by (20th-cent.) poets, and as nonsense when they believe them to have been computer-generated.

I have left out, and I may not even have properly described the identificatory features of this one line, but you get the idea. We manage to identify metaphors in many unrelated ways, and what matters here is that (h) is indeed a metaphor.

The second question is not completely independent from the first. In addition to considerations of the previous paragraph, it may well be that our judgement that (h) is a metaphor is partly the result of, and partly responsible for, our casting around for a subject *other than* the event with the ball. The poem from which the line is taken is autobiographical, or at least purports to be. Even without the surrounding text, it seems clear enough that (h) is likely to be about the course of some individual's life, though the sentence itself makes no explicit reference to that life. By my lights, what is asserted is that the event qualifies that life, that we understand something about an individual life from it. To be sure, it is difficult to use literal language to say what we thereby understand, though this is as it should be given the impossibility of paraphrase in general, and the vividness of this metaphor in particular. But there are elucidations and rationalizations that could help orient someone who is completely perplexed by (h). One might begin, in an elucidatory mode, by saying something about the way seemingly small events in our earliest childhood lay the foundations for, and are mirrored in, the seemingly more significant events of adulthood. And then, speculating about the author, one might add in the rationalizing mode that when Thomas wrote this, he was obviously at that stage in life when he felt he ought to, but couldn't, shake off the psychological effects of childhood experience. (This is speculation, solely for the purposes of the exposition.)

In sum, we have in (h) a sentence which is obviously enough metaphorical—in some sense, the whole sentence is—but which does not explicitly include the subject of the metaphor. And even more than in some previous examples, while it would be absurd to attempt to twist (h) into subject-predicate form, there is still an underlying hint of 'this-is-that' about it. On a semantic descent account all of this can be accommodated. The sentence as a whole, via semantic descent, gives us an object: the event in the playground. The sense of 'this-is-that' comes directly from the fact that this object of metaphor serves as a qualifier of the implicit subject of the metaphor, an individual's life.

One sentence from the original list remains:(a) Out of the crooked timber from which men are made, nothing straight can ever be built.I saved it for last, not because it introduces any special further problems, but simply because it seems a particularly nice case for the semantic descent account, as well as being a powerful metaphor. As in other cases, there are trade-offs in the various ways of marking semantic descent. Here are two:(a1) ↓ ↓Out of the crooked timber from which ↓men↓ are made, nothing straight can ever be built↓ ↓.(a2) ↓Out of the crooked timber↓ from which men are made, ↓nothing straight can ever be built↓ ↓.

The oscillation is set up by the fact that we can be either literally speaking about what men are made from, or, more simply, just of men in relation to the fact that no genuinely straight thing can be made from crooked timber. The effect of this oscillation seems to me both like that of ‘great’ in the Tolstoy sentence, and like the hedging necessary to smooth over awkward juxtapositions found in several of the other sentences. It serves to focus attention both on what human beings are made from, and on the consequences of this for human nature. And we are invited to understand these things without ever having had to juxtapose directly—by subject-predicate confrontation—men and crooked timber, or men's devious natures and the products a carpenter can make out of crooked timber.

I have now discussed all of the examples on the original list, as well as a sentence about Tolstoy that was not on that list. The next main task of this chapter is to consider issues of metaphor complexity *other than the merely syntactic*. However, just before moving on, I should like to add two vital observations that arise directly from the discussion of syntactic complexity.

4.2.5. Sharing properties

Having run through a varied set of examples, I have shown how the semantic descent account deals with the challenges of explaining both syntactic complexity and the underlying ‘this-is-that’ nature of metaphor. However, even though my examples take us substantially beyond the subject-predicate form, they are still restrictive. This comes in part from my having had to focus on sentence-sized metaphors—what else, since I have been considering syntactically more complex *sentences*—and in part from the fact that my examples needed to be self-contained enough to be intelligible outside of their larger linguistic contexts. Later in the chapter, when I come to consider other kinds of complexity, I will say something about the extension of metaphor beyond the unit of the sentence. Before moving on, I should like to use the just completed survey of examples to address an issue which has surfaced more than once, and will surface again. I have in mind here the widespread appeal to the idea that metaphor involves a sharing of properties.

The example central to the first three chapters certainly lends some support to this appeal. In one way or other, many authors assert or assume that what Romeo is doing in uttering (R) is saying that Juliet and the sun share one or more properties.¹³⁰ And this seems plausible given that (R) confronts us with two things, Juliet and the sun, brought together in a way which might suggest comparison, and therefore the potential sharing of properties. In previous chapters, I have argued that this idea that properties are shared is not simply mistaken—not, at least, if it is handled carefully—but that it is hopeless as an explanation of what Romeo is doing in uttering (R), and hence hopeless as an explanation of metaphor in general. In the next chapter, I will have occasion to bring these arguments together, but here I want simply to note how lame the idea of property-sharing is when one thinks of more syntactically complex examples.

¹³⁰ Sometimes talk of property-sharing is explicit in an account, but sometimes it is only implicit, as when an author simply speaks of comparing things, but then spells out comparison in terms of property-sharing. For present purposes, accuracy of attribution to this or that author is not important.

In (e), an argument is characterized by the swerving and sudden stopping of a vehicle barely under control. Does it make sense here to look for properties of the latter that might be shared by the argument, and which would thus explain the metaphor? It certainly doesn't seem to. What does the actual work in the metaphor is the swerving and sudden stopping of a vehicle barely under control, not properties of its swerving and sudden stopping. To be sure, once one has established that the vehicle's movements characterize the argument, one might set out to list features of the movement which might correspond to features of the argument. As I have argued before, this quest is less straightforward than is usually recognized: for example, the predicates 'swerving' and 'lack of control' when used of a vehicle do not pick out the same properties when applied to an argument; the words in the predicates may be the same, but the properties certainly are not. Still, this is not the present difficulty. What I am stressing here is that, unlike simple subject-predicate metaphors, the ingredients of more complex metaphors do not encourage us even to begin the search for common properties, at least until we have already established what characterizes what in the first place. And, of course, 'characterizes' is merely another way of expressing the non-linguistic predicational function that I have labelled 'qualification'.

The sharing of properties view looks lame when thought of in respect of most of the examples in the original list. Even the relatively simple (f), where it is an event of shining that is used to characterize Juliet's entry, doesn't fit well into the property-sharing story. It is the event of, say, the sun's shining that does the characterizing work, not properties of the sun's shining. However, the more complex phrasal metaphors make the point even more emphatically, and none more so than the subtle (a), or the sentence-sized example (h).

In (a) one might insist that there is an underlying juxtaposition of men and timber, and one might think it reasonable to look for properties of the timber that are somehow shared by men. But aside from the fact that there are almost certainly no interesting properties of the one that are properties of the other, (a) doesn't invite this way of talking. What is said is that only unstraight things can be fabricated from crooked timber. This fact of carpentry, not properties of it, tells us something about human development.

The same moral can be drawn from (h), though here it is at its most forceful. First off, we have to ferret about to find a subject for the metaphor—an individual human life—but, once found, it is not *properties of* the 5 year old throwing up a ball in the playground, but that event itself which does the characterization. As before, one can speak of properties of the event—that it was done by a 5 year old, that it is not yet completed, etc. which are shared by that life. Remember that I do not claim that it is mistaken to describe a metaphor subject as sharing properties with the object of metaphor. But this talk of property-sharing is *consequent* on our having identified the object of metaphor—the qualifying object—in all its specificity, and it is the use of this object which figures in the account of metaphor. Property-sharing talk is at best something which comes later, and might figure in what I have called 'elucidation'. It certainly doesn't explain the metaphor in the first place.

4.2.6. Attunement again

In Chapter 3, I struggled to make out a case for the truth aptness of metaphor assertion. The obstacle to be overcome was the tendency to think that, when objects function as qualifiers or proto-predicates, they are insufficiently determinate in that role to give truth-conditional content to putative assertions. To revert to the original example, when Romeo says that Juliet is the sun, and this is treated as Romeo's having asserted: (R') Juliet is ↓the sun↓, the sun itself fills out the role of the predicate. But the worry was that the sun, used in this way, is not sufficiently constrained to make (R') a report of an assertion with a determinate content. I won't here repeat the arguments that I used to allay this worry. In outline, they involved, on the one hand, stressing that even linguistic predicates do not have the determinacy they are often imagined to have, and, on the other, that our knowledge of both context and the 'cultural' significances of objects attunes us to their qualifying roles. The point I want to emphasize here is that the issue of determinacy is much less of a worry in respect of more complex metaphors. I had suggested as much at the end of Chapter 3, but was constrained there by my wanting to present the semantic descent account without embellishment.

Having already worked through my list of sentences in perhaps too much detail, let me just choose one to make the point. When Kant's metaphor is marked for descent as: (a1) ↓ ↓ Out of the crooked timber from which ↓ men ↓ are made nothing straight can ever be built ↓ ↓, we are invited to make the descent, first to a fact of life for carpenters, and then to a use of this fact as a qualifier of human beings and their development. (There is no need for the present point to consider alternative ways of marking (a) for descent.) Now, I have no doubt that there will be subtly different ways in which this proto-predicational act will be taken, but surely there is nothing indeterminate about the object which serves to effect this act. Unlike the case with 'the sun', the complexity of what lies between the down-arrows guarantees that anyone who understands (a) is in no doubt about shape of the predicative role of the object called on, even though elucidation and further commentary are both appropriate.

4.3. Mixed and Extended Metaphor

As part of the run-up up to considering non-syntactic complexity—features such as vividness and richness—I begin with some comments on mixed metaphor, a topic often neglected in the literature.¹³¹ There are three reasons for doing so: it is an intrinsically interesting and even puzzling topic; what is said about it will contribute to the more substantial issues discussed in the sections which follow; and, finally, the semantic descent account offers the best hope of dealing with it.

¹³¹ That said, there is valuable material on this topic in both Cooper 1986 and White 1996.

Any reasonable course of instruction in writing advises students to avoid mixing metaphors. But though the advice is well-known, the reasons for it are at best unclear. Indeed, on some philosophical accounts of metaphor, there would seem to be no disadvantage, and possibly some positive benefit, in such mixing. If a metaphor's meaning results from some alchemical process involving the literal meanings, connotations, and commonplaces or properties associated with words, then mixing of metaphors might actually speed up the process, or make its outcome more determinate. Or, at the very least, it might be felt that, as the saying goes, the more the merrier.

Of course, as is so often the case with literary advice, the stricture against mixing metaphors is not universally apt: it is not all that difficult to find examples where apparently mixing metaphors works well enough. So, perhaps the fact that certain philosophical accounts give no principled reason for avoidance is no weakness on their part. Perhaps the stricture against mixing metaphors is little more than a stylistic prejudice, without any sound semantic basis.

This line of defence is desperate. Mixed metaphors are often, if not always, a bad thing in a pretty straightforward way: they can sow confusion where none was meant to be offered for sale. (See what I mean.) Any account of metaphor should be able to explain this. Indeed, what would be best would be to have an account of metaphor that gives clear reasons for the ban, while at the same time explaining why exceptions sometimes work. My suggestion is that the semantic descent account does just this.

First let's consider some obviously unacceptable cases of mixed metaphor. In advance of any more theoretical work, I have to rely on our intuitions that these are unacceptable, but the examples below will not I think strain those intuitions:

- (1) Our society has a dog-eat-dog pecking order.
- (2) The ship bowed to the waves by stepping gently over each one in turn.
- (3) They chafe at our impotence if, facing the would-be knave, we try to drum him into the ranks of the virtuous but find nothing on which to get a purchase.¹³²

Each of these is a case of blatantly mixed metaphor, but there is no need here to discuss all three, so I shall let the middle one serve as the main culprit for expository purposes.

Notice first that the two metaphors mixed in (2)—one which has the ship as some sort of courtier bowing to the waves, and one which has it as a careful pedestrian stepping over them—both occur in the same sentence. So far, I have restricted myself to considering features of metaphor that show up in sentence-sized units;

¹³² The first of these is taken from a *Guardian* newspaper survey of how badly children performed on various linguistic tasks (it was one of many mixed metaphors that occurred in students' essays). The final example is from a contemporary philosopher, and the middle one is based on something I recently heard on the radio. The discussion will be mostly conducted in terms of the latter, but I included the other two to remind you that hopelessly mixing metaphors is ubiquitous and occurs at all levels of sophistication in writing.

a discussion of metaphor in larger units of discourse will come later. But when it comes to mixed metaphor, this restriction is actually quite important. It is often unclear whether the use of several different metaphors in a text longer than a sentence is ever a case of mixing. Whether it is or not depends on various internal features of the discourse and its setting. Moreover, even what is grammatically a single sentence can contain more than one metaphor without that counting as a case of mixed metaphor. However, when two metaphors occur as they do in (2), then we have a clear instance of mixed metaphor, and one which is certainly not acceptable. The question is why is this so?

On my account, any intelligibility-conferring interpretation of (2) requires us to move from certain of the words in the sentence to objects, events, or states of affairs in the world which then serve a predicative, or more accurately a qualifying, function. Initially, we reach for some event indicated by the word 'bowing' and then use it to qualify the ship's relation to the waves. But no sooner do we set out to do this than we are brought up short by having also to think of some other event, one indicated by the phrase 'stepping gently over'. Why are we brought up short? Because of the sheer impossibility of finding some event which is both a bowing and a stepping gently over. Unlike views which demand the creation of a new metaphorical meaning for the words in the sentence, mine requires us to employ the absolutely standard literal meanings in the search for some real-world item. While it is certainly possible to think that some new predicate meaning could be generated from the meanings of the words in (2)—some non literal characterization of the ship's motion—there is simply no way in which those words could lead us to a single event that could then be employed predicatively. Mixed metaphors are thus to be avoided simply because they are uninterpretable.

But not always. Consider this sentence: (4) His life was a setting sun, burning the last remnants of his ambition. Here we have a case in which metaphors are undoubtedly mixed: someone's life is compared both to a setting sun and to a fire whose fuel is nearly exhausted. But, unlike (2), which is unacceptable, even comic, (4) somehow seems alright. It might not be particularly original or striking, but it does not bring us up short. Why is this? There is an intelligible interplay, not simply between the words used in each metaphor, but between the objects we access by semantic descent in each case. The first metaphor invites us to think of the sun's setting as the natural predicate of a certain life; in the second a burning down of a fire is used to the same purpose. These are of course quite different items: the sun is not, in setting, burning the last of its fuel, and a fire does not burn down by disappearing behind the horizon. However, it is not unreasonable to think that we could have used the object indicated in the second metaphor as a way of conveying information about the first. Somewhat tritely, but intelligibly, we can say *metaphorically* that the setting sun is a fire burning the last remnants of its fuel. No less plausibly, we can use the object indicated in the first metaphor to convey information about the object in the second. A fire burning down could be described as a setting sun. It is this metaphorical

interplay between the two proto-predicates that saves (4) from the blank unintelligibility of (2). In (4) we have a case of mixed metaphor that works, or, if you find my example somewhat forced, is at least less jarring than (2).

The conventional advice to avoid mixed metaphor was surely framed with examples like (2) in view, and, on the semantic descent account, the reason for this is clear: to mix metaphors in this way is to make what you write uninterpretable, or near enough. With examples like (4), the aptness of the advice is less clear cut. Strictly speaking, semantic descent in (4) leads to the same kind of interpretative cul-de-sac as (2), but there is a difference: the different objects got by descent could be used as qualifiers of each other. When this happens, one might find that the mixed metaphor works perfectly well. Moreover, while this possibility helps us understand the success or failure of mixed metaphors, it has a much wider application. For it can happen that two different words or phrases might lead by semantic descent, not to two different objects, as in mixed metaphor, but to the *same one*. This possibility can give us insight into the ways metaphors become extended over larger stretches of discourse than individual sentences.

To get an idea of what I mean, consider one of the examples from section 4.2:(5) When questioned, he offered his usual soap-bubble reason for what he had done. As noted, the phrase ‘soap-bubble reason’ bears the burden of metaphor in (5), and this phrase functions by calling on an object—a soap-bubble—which then serves as an adjectival qualifier of ‘reason’. Imagine this sentence followed by:(6) But this was pricked by the detective asking about the telephone call to the bank manager, and it burst completely when the bank official denied that any such call took place. This continuation shows how a metaphor can reverberate in what one could call (grandly, in this case) the larger discourse. Whilst neither particularly vivid nor profound, the metaphor reverberating its way through (5)–(6) is far from dead, and it is a useful reminder of the kind of transparency and currency that metaphors often have. In any case, it will serve well enough as an example for showing how to deal with extended metaphors.¹³³

My account offers a simple, and I think natural, explanation of what happens when a metaphor is extended. What makes extension possible in (5) is that the object of metaphor remains available for further employment. The reason for action, as characterized in (5), encourages contributions to the continuing discourse which are soap-bubble-related. And, most importantly, these contributions are each further instances of semantic descent and qualification. Thus, the detective's question is qualified by an event in which someone prods a balloon or bubble with a sharpish

¹³³ The label ‘extended metaphor’ is borrowed from White 1996: 144–54, but discussion of his treatment of the phenomenon won't come until Ch. 5. I should also say here that I don't think we need to worry about the fact that a metaphor can be extended within a sentence or within a larger text made up of several sentences. It is how extension works that counts, not how far it reaches.

instrument, though, given (5), it is no surprise that we take the relevant object to be a soap-bubble; and the rather more forceful prodding that leads to the bursting of this same soap-bubble is the event which characterizes the actual denouement of the detective's line of questioning, the exposure of the original reason as worthless.

There can be an interesting overlap between the phenomenon of extended metaphor that we see in (5)–(6), and a certain kind of commentary on metaphor. Recall that one of the ways in which we comment on metaphors is by saying more about what by my lights is the qualifying object. Thus, faced with someone who is having trouble with Romeo's (R), we might well come up with one or all of: The sun casts the light which makes things visible, The sun warms us, The sun is necessary to make things grow. These kinds of comment were described earlier as *elucidations* of the metaphor. However, in the right context, these or closely related remarks, might be understood more as extensions than elucidations. Thus, if not in the spirit of helping someone with interpretation, one responded to (R) by: One needs protection from the sun, this should certainly count as an extension of the metaphor, though its relationship to the metaphor in (R) is more complex than the relationship of (6) to (5). One could say that it is an extension, though not a continuation, of the metaphor in (R).

Much more could be said about extension, elucidation, and the overlap between them, by drafting in richer and more interesting examples.¹³⁴ However, I don't think that anything would be added to (or subtracted from) the outline story just told.

4.4. Dead Metaphor

As noted at the beginning of the chapter, philosophers often appeal to examples of metaphor which are implausibly simple in syntactic structure and are neither vivid, inspirational, creative, nor, to use a popular expression, 'high-octane'. Thus, talk of human beings as lions, foxes, rocks, blocks of ice, and pigs bulks large in philosophical writing about metaphor. Having already dealt with syntactical complexity, my aim in this section is to begin addressing the other kind of complexity in metaphor, complexity we might best summarize as the *richness* of metaphor.¹³⁵

¹³⁴ White (1996: 144–54) offers two Shakespearean examples of extended metaphor: the first 'a case of a single continuous passage that develops a metaphor through two or more sentences'; and a second in which 'the extension is provided by a different speaker, who puts it forward with an intention radically at odds with that of the original speaker'. Coincidentally, these illustrate my examples of pure extension and extension by elucidation.

¹³⁵ For reasons of expositional simplicity, I shall conduct the discussion of dead metaphor almost entirely in terms of the usual subject-predicate examples. Having dealt in detail with syntactic complexity, I do not think that doing this opens me to the charges of oversimplification, and it will make the discussion easier to follow.

Not interesting in themselves, the tired examples noted above nonetheless share a feature which promises more: they are all cases of what is usually, and sometimes pejoratively, called ‘dead metaphor’. The topic of dead metaphor is more interesting than its instances, because if we could give a plausible account of it, then by contrast we might be able to gain some insight into the phenomenon of live or vivid metaphor. This is certainly something I hope to convince you of in what follows. Additionally, we can use it as yet another yardstick against which to judge accounts of metaphor: any account worth considering should be able to explain the phenomenon of dead metaphor.

That said, there is an initial problem with this strategy. The expression ‘dead metaphor’—itself metaphorical—can be understood in at least two ways. On the one hand, a dead metaphor may be like a dead issue or a dead parrot; dead issues are not issues, dead parrots, as we all know, are not parrots. On this construal a dead metaphor is simply not a metaphor. On the other hand, a dead metaphor may be more like a dead key on a piano; dead keys are still keys, albeit weak or dull, and so perhaps a dead metaphor, even if it lacks vivacity, is metaphor nonetheless. (Another example: when you have overdone physical exercise, you might describe yourself as ‘dead’. But of course you don’t mean that you have ceased to be—that, like the ex-parrot, you are no longer a functioning human being—what you really mean is that you are tired. So, perhaps a dead metaphor is a tired one that might come to life, so to speak, after some sort of rest.)

It would be good if we could decide precisely how to unpack the metaphor in ‘dead metaphor’ *before* we consider the relationship of this phenomenon to accounts of metaphor (my own included). However, as we have precious little leverage on what constitutes a dead metaphor, this is not really an option. So, making the usual virtue of necessity, I shall consider the notion of dead metaphor against the background of types of account classified in Chapter 1.

4.4.1. How dead is dead?

Reverting to the original example, the problem for any philosophical account of metaphor is that of making intelligible Romeo’s utterance of: (R) Juliet is the sun. The Content Sufficient theorist, as the label implies, finds there to be some meaning or content which allows us to understand Romeo’s linguistic act. Thus, to take Black’s interactionist view (or a simplified version of it): the predicate in (R) has associated with it various commonplaces out of which, and in interaction with Juliet as subject, we construct a second meaning or content which renders the act of uttering (R) intelligible. In contrast, theorists like Searle and Davidson treat the words in (R) as having only a single meaning, one that would lead us to judge it as palpably false. So Content Insufficient theorists must say something more about what Romeo is doing. It is over this further thing that Searle and Davidson part company. Searle’s view is that we find some substitute for what the words of (R) mean, some speaker meaning which Romeo uses (R) to get across, whereas Davidson denies that there is any such further message. He would have us understand Romeo as doing

something else, something he characterizes as getting us to see Juliet as the sun, or as putting an image in front of us. In the present context, the question is: what consequences do each of these alternatives have for our understanding dead metaphor?

Answering this will be easier if we have an example in front of us. Suppose that Romeo had also said: (7) Capulet is a fox. This is undisputedly a dead metaphor (possible anachronism aside), but is it nonetheless a metaphor? A Content Sufficient view like Black's might deal with this question as follows. At some point in its history, (7) was a live metaphor, one requiring a hearer to put in some effort to work out the commonplaces associated with 'fox', which, in interaction with the fact that the subject of (7) is human, produced a metaphorical meaning for this word. However, given that one can now find this meaning in any decent dictionary, such effort is no longer required, and the metaphor in (7) is dead. One could in fact say that 'fox' has become a polyseme: a word with several different though related meanings. Nonetheless it is still true that, in the absence of knowledge of current English usage, someone could come across this use of 'fox' for the first time, and work out its metaphorical meaning. It is for this reason that (7), while dead, still counts as a metaphor for Black.

As a treatment of *live* metaphors, and especially those whose syntax is more complicated than (7), Black's account is beset with problems, many of which have already been discussed.¹³⁶ But here the only issue is how dead metaphor is handled, and there seems a clear moral we can draw. Not just Black's, but Content Sufficient accounts in general, allow metaphors to be dead on their feet, so to speak, without losing their status as metaphors. Essential to this are two ingredients: the possibility of attributing metaphorical meanings to words or expressions, and the fact that effort is needed to derive metaphorical meanings from the original literal meanings of words and expressions.

In contrast, there is an equally general reason why Content Insufficient accounts must regard dead metaphors as *ex*-metaphors. Take Searle's version: he insists that in genuinely live metaphor the speaker uses words with their literal content to put across some further message. Now deciphering this message obviously requires work on the part of the hearer, and the need for such work is something Searle's view shares with Black's. But when, as with 'fox', a meaning is available in the dictionary, not only is no extra work needed to arrive at it, but the very idea of a further or alternative message lapses. And without even the possibility of speaker meaning, Searle cannot regard 'fox' as metaphorical at all. In Davidson's version, there is no alternative or further message, but as far as the deadness of a dead metaphor is concerned, he ends up even more quickly in the same place as Searle.

¹³⁶ One of the troubles with views like Black's is that it has difficulty in acknowledging properly the fact that the literal meaning of a word is at actually at work in the metaphorical meaning, and is not simply, as Moran 1997: 253 puts it, 'a ladder [to metaphorical meaning] that is kicked away'. Perhaps then it is unsurprising, given the kinds of examples Black and others use, that such views do better with dead than with live metaphor.

Since Davidson insists that there is no message in a metaphor that explains what the speaker is up to, the fact that a dictionary reveals just such a message rules out (7) as a metaphor.¹³⁷

The upshot of the last couple of paragraphs is clear, but by itself inconclusive: Content Sufficient accounts seem committed to regarding dead metaphors as tired; Content Insufficient accounts to regarding them as no longer metaphorical at all. But who is right about dead metaphor? As you might expect, the evidence is not absolutely clear cut, but it does lean heavily in favour of the idea that dead metaphors are metaphors. There are three sources of evidence, though they overlap.

4.4.2. Sifting the evidence

The first, and in some ways most significant, piece of evidence has a certain irony to it. I have expressed my disappointment that philosophers often use what are dead metaphors in their discussions. But the plain fact is that these are used as, and tend to be accepted as, examples of metaphor. To be sure, these are often said to be low on the octane scale, suitable only for expository purposes, etc. But the very fact that a discussion of metaphor can be conducted using them suggests very strongly that they are live enough to count as metaphors, in spite of their being dead enough to figure in dictionaries.

The second kind of evidence is more subtle. It can happen that a dead metaphor can be brought to life by either inadvertent carelessness, or quite deliberately. But of course this only makes sense if the thing in question is not really an *ex*-metaphor. The absurdist humour in the Monty Python parrot sketch is most evident when the pet-shop owner tries to convince the customer in various ways that the dead parrot might only be resting. But there is nothing absurdist about the following ways in which a dead metaphor can be revived.

Mixed metaphor is one such way, and one of the examples in the last section illustrates this nicely. If someone complains about modern life by saying: It's a dog-eat-dog world out there, there can be little doubt that the metaphor here is as dead as can be. But when a careless student writes: Our society has a dog-eat-dog pecking order, the canine metaphor in 'dog-eat-dog' comes to life, if only to protest at being combined with an equally dead metaphor more appropriate to fowl.

A more delicate example discussed by Cooper (1986: 128–9) comes originally from Fowler. It is clear enough that our talk of 'sifting evidence' is as good an example of dead metaphor as any so far considered. But when we read: All the evidence must be carefully sifted with acid tests,

¹³⁷ Cooper (1986) reaches this very conclusion about Davidson and Searle, but he doesn't consider the phenomenon of dead metaphor from the perspective of what I have labelled Content Sufficiency. I am in a general way indebted to Cooper's detailed discussion of this phenomenon—especially in the few paragraphs that follow this one—even though I draw a radically different conclusion from it.

the very fact that ‘metaphors’ are mixed seems to bring these dead metaphors to life. Cooper suggests that Fowler's example is not conclusive—that the fact that ‘sift’ and ‘acid test’ were *once* metaphors might make the above only stylistically awkward, without threatening its intelligibility. However, this evasive action doesn't seem to work as well in the canine/fowl example, and in any case there is a whole range of further examples which make the ‘merely stylistic’ option unattractive.

The two examples above might be thought of as Frankenstein cases: their authors brought to life something which proved self-destructive. This is what you would expect if inadvertently mixed metaphors become lively enough to render incoherent the sentence in which they occur. But, as we saw in the previous section, there are ways of mixing metaphors *deliberately*, even in a single sentence, that have no such self-destructive effects. In these cases, there is a kind of resonance created by one metaphor interacting with another that enhances the overall effect without in the least threatening any kind of incoherence. Given this, it should in theory be possible to find cases in which the beneficial mixture is of dead not live metaphors. Nor is this only a theoretical possibility, as witness: ‘Tom is a snake posing as a fox,’ in which we have the mixing of two dead metaphors that, through our efforts to cope with the implausibility of the mixture, brings each to life. (This works even better when the mixture of metaphors occurs across sentences in a larger discourse. There one can achieve all sorts of the kind of resonant effects described in the last section, even though the metaphors used are, shall we say, tired.)

The third kind of evidence arises from my earlier discussion of extended metaphor. It is not difficult to see how we can revive virtually any dead metaphor by constructing appropriately extended contexts or continuations. Thus, ‘mouth’ in: ‘The ships entered the mouth of the river,’ is dead, but it comes to life again in: ‘The lips of the river's mouth parted with the rising tide to let the ships enter.’

Summing up: three sorts of evidence point strongly in favour of the idea that a dead metaphor is nonetheless a metaphor. This is not good news for Content Insufficient accounts, but my interest here is not so much in general arguments against these views, as it is in the phenomenon of dead metaphor itself.

Cooper (1986) adduces a fourth kind of evidence which he finds compelling. It is evidence which comes from the so-called ‘cognitive’ account of metaphor championed by Lakoff and Johnson (mostly in 1980, but further explored in Lakoff and Turner 1989). It will be useful here to say why (*pace* Cooper) I do not regard Lakoff and Johnson's discussion as contributing anything to the present issue, over and above what has been said already, nor indeed as being all that helpful in respect of metaphor in general. Doing so will lead naturally to my own account of dead metaphor.

4.4.3. The cognitive account

Lakoff and Johnson claim that the phenomenon of metaphor has less to do with language than with basic processes of thought. Thus, perhaps surprisingly, they insist that we should not regard metaphor as this or that kind of *linguistic* construction, but rather as a kind of thought that underlies and generates what would otherwise be isolated uses of language. The view is too familiar to need much exposition, so let one example serve. We speak about time in all sorts of apparently metaphorical ways: we say that we invest time, spend time, waste time, etc. Each of these could be regarded as individual metaphors in need of explanation, but Lakoff and Johnson suggest instead that the real metaphor is the underlying *thought*: TIME IS MONEY. It is this thought which unifies the scattered ‘metaphors’ we find irresistible when speaking about time. (They express the thought in upper case letters as a way of indicating that it is a conceptual device, and not itself a linguistic metaphor.)

Now it has been pointed out that many of the examples Lakoff and Johnson use to illustrate their view are in fact dead metaphors, and that this makes their view unimpressive as an account of the phenomenon itself. For reasons which will become clear, I don't think that this criticism is particularly damning, and anyway, as we have already seen, the use of dead metaphors is par for the course in many discussions. Besides, dead metaphor happens to be what is under discussion here, and Cooper's interest in the Lakoff and Johnson cognitive account springs from its supposed contribution to that discussion.

Having noted that Davidson and Searle (and all Content Insufficient theorists) must regard dead metaphor as ex-metaphor, Cooper asks the obvious next question: what have these former metaphors become? The answer which he regards as most plausible is that a dead metaphor—he prefers at this point to speak non-pejoratively of ‘established’ metaphor—leads to the inauguration of additional literal meanings, to polysemes. However, though most plausible, he does not find this plausible enough, because of the systematicity and generative power inherent in the phenomena revealed by the cognitive account. In respect of systematicity, Cooper notes: ‘[I]t is the exception, rather than the rule, for established metaphorical expressions to have become established singly. More typically, it is an expression along with many other related expressions which, en bloc as it were, develop a new usage outside of the parent domain’ (1986: 130–1). As Lakoff and Johnson insisted, established metaphors such as ‘spend time’, ‘invest time’, ‘waste time’, and others in this same vein, seem systematically related. Moreover, Cooper insists that this aspect of systematicity cannot be mimicked by the polyseme option: One thing linguists sometimes mean by referring to a linguistic practice as ‘systematic’ is that it has a *generative* power; that it gives rise to novel utterances which are readily interpreted only because people are acquainted with the practice in question. Now polysemes, in the rare cases where they come in groups, are not generative. (Cooper 1986: 133)

Here too he goes on to describe typical examples from Lakoff and Johnson's account, examples in which we apparently lean on some underlying ‘metaphor’ such as TIME IS MONEY in both our comprehension and, most especially, in our

finding there to be extended batteries of linguistic metaphors related to each other. Such extension cannot be mimicked in any principled way by the polyseme view, so Cooper finally rejects it: It would surely be bizarre if we were to approach such central questions about metaphor as ‘Why do we speak metaphorically?’, or ‘Can metaphor provide a distinctive kind of understanding?’, without even taking into consideration the batteries of systematic established talk that I have been referring to and illustrating. Yet on the polysemy view, we should not be entitled to let them into our considerations, since their metaphoricality is a mistake. (Cooper 1986: 135–6)

Cooper aligns himself with the Davidsonian account of metaphor, so for him this conclusion is seriously troubling. On the one hand, he is committed to denying that dead metaphors are metaphors; but, on the other, he thinks that the systematicity and generativeness typical of dead metaphor rule out the only remotely plausible account available to Davidson of what a dead metaphor becomes. In this tough spot, Cooper chooses a heroic option: he draws a distinction between ‘speaking *metaphorically*’ and ‘uttering *a metaphor*’, and then insists that in using a dead metaphor one is not uttering a metaphor, though one is speaking metaphorically. As he himself admits, this is not an easy distinction, nor does it fall in with the ways we use these words. Still, not myself being in the position of having to defend the Davidsonian conception of dead metaphor, there is no need for me to look more closely at Cooper's conclusion. What matters here is the appeal to the Lakoff and Johnson cognitive account, and in particular, the question of whether this account adds anything to the evidence already given for the revivability of dead metaphors.

I think not. Nor, as suggested earlier, does the cognitive account of metaphor add anything to the account that I have been developing in this book. The reasons for these conclusions are connected, and both will eventually lead to my own account of dead metaphor.

4.4.4. Semantic descent and the cognitive account

The semantic descent account is clearly enough Content Sufficient, even though it differs radically from the other views in this category. In the present context, this means that it shares with other Content Sufficient accounts the resources necessary for treating dead metaphor as metaphors. As we saw earlier, the key to this lies in the joint possibility of there being metaphorical meanings in addition to literal ones, *and* of there being some work required to explain utterances that call on those meanings. One can then treat a dead metaphor as an expression whose metaphorical meaning can be obtained *without* this effort, since, for example, it may well figure in a dictionary, but which, in the right circumstances, could still be so obtained.¹³⁸

¹³⁸ As I have stressed from the beginning, my account is a philosophical one: it aims to show how to treat metaphor within a philosophical account of meaning. It is not therefore an account of the psychological processes which speakers and hearers engage in, though, if handled carefully, it does have consequences for psychological theory. Still, in the discussion of dead metaphor, it is difficult to avoid psychological talk. Clearly, the very phenomenon of dead metaphor is itself entangled with the perceptions of actual speakers and hearers, so it is difficult to avoid giving the impression that the whole of the semantic descent account is itself a psychological one. Still the basic question is one of which resources we need to call on to account for the intelligibility of this or that utterance, and admittedly this depends to some extent on how those utterances strike speakers and hearers. Thus, the resources we need to explain talk of rivers' mouths in the ‘dead’ context will be different from those needed in a revivifying one. I have often put the point in terms of the efforts of speakers and hearers, but this psychological talk can be misleading. It would be possible each time to transpose apparently psychological talk about how certain utterances strike speakers into talk of what is necessary to account for the intelligibility of linguistic acts effected in these utterances. But, as already noted, such repeated transposition would lead to unacceptable awkwardness in formulation. That said, the issue of psychology and theories of meaning is delicate, and is not one systematically addressed in the book.

Unlike other Content Sufficient accounts, the semantic descent view postulates two quite specific and separable kinds of work that figure in the generation of its rather special kind of metaphorical meaning: there is the descent from ordinary linguistic expressions to objects; and there is the use of these objects as qualifiers or ‘proto-predicates’, as I sometimes call them. The fact that there are these two components is of special importance to the phenomenon of dead metaphor.

Throughout, I have been very careful to insist that metaphor requires both of these; qualification on its own, real as the examples in Chapter 2 showed it to be, is simply not metaphor. From this perspective, the cognitive account can be seen as a rather confused attempt to treat something rather like qualification as though it were itself metaphor. It is of course only something like qualification because none of the background discussion of this notion, or anything close to it, figures in the cognitive account. Still, what Lakoff and Johnson regard as ‘conceptual’ devices lying behind a whole range of linguistic metaphors—‘formulae’ like TIME IS MONEY—are by my lights a kind of qualificational claim: money (by which is meant a whole range of determinate monetary transactions) qualifies time. To be sure, formulae like TIME IS MONEY or LOVE IS A JOURNEY lack the specificity of the qualificational examples discussed in Chapter 2, but they can be seen as gestures in that direction. This will be clearer if I say (briefly) how the cognitive accounts looks from my perspective.

People say such things as: (8) Professor X is not spending his time wisely. When they do, we could imagine there to be semantic descent to some object, in this case a course of action in which money is not spent wisely, and this object will then qualify the subject of the metaphor, in this case Professor X's course of action. Many such metaphors are possible, and they cluster around a theme in the qualifications that could be called on in each case. This qualificational theme can be summarized (without the capital letters) as: time is money. But the latter is not itself a metaphor. Thus, far from thinking that some cognitive device is the *real* metaphor grounding a range of common ways of speaking, that device simply indicates that the relevant range of metaphors potentially makes use of a recognizably similar kind of qualifying object.¹³⁹

¹³⁹ Fogelin (1988: 85–6) and White (1996: 300–1) both express reservations about the cognitive account different in detail from mine—something not surprising since my objections grow out of a rather different account of metaphor from theirs—but not all that different in general thrust.

Note that I have been careful not to claim that when we use metaphors like (8) above we actually make the effort to descend semantically and work out which object does what kind of qualifying. Virtually all of Lakoff and Johnson's examples involve metaphors that are dead, and from my perspective what this means is that we do not actually need to call on descent or qualification in dealing with them. But, in the right circumstances, we might be forced to do so. Thus, when Shakespeare writes (*Troilus and Cressida*, 3. 3. 139–42):
 Time hath, my Lord, A wallet at his back,
 wherein he puts Alms for oblivion,
 a great-sized monster Of ingratitude,
 the dead metaphors that depend in general on the range of ways in which money qualifies time comes to life in a startlingly vivid way. By my lights, our attempts to understand Shakespeare here involve our first imagining an individual wearing on his back a money bag of some sort in which he keeps sums intended for an, as it happens, ungrateful recipient, and then taking this scene to qualify the way in which we carefully meter our time in the attempt somehow to slow or prevent our inevitable death. (Tempting as it is, I shall not say more about the cleverness of this image except to note how the semantic descent is guided by using 'time' and 'oblivion' as if they were proper names of items which figure in the qualification.)

The above observations suggest then that, on the one hand, the cognitive account does not in fact require us to change our view of what a metaphor is. The formulae underlying groups of linguistic metaphors are not themselves so-called 'cognitive metaphors' (or 'metaphors in thought'); they are not metaphors at all. On the other hand, while these formulae can be understood as pointing to a commonality in the qualifiers which figure in genuinely linguistic metaphors, there is no reason to think that we appeal to descent and qualification in our understanding of them. Metaphors involving time and money, for example, are dead, and this means that they are available to speakers of English without their having to put in the effort required of live metaphors. That said, it is always possible to breathe life into them by any one of the means touched on earlier. We do not pause over the dead metaphor in 'mouth of a river'—no descent to the object and no use of the object to qualify the river—but we can be forced to reflect on these metaphor processes when this expression is inappropriately mixed with other metaphors (dead or alive), or when it is used in contexts that force us to appeal to these processes. So, the cognitive account is not radically alternative to the semantic descent account and certain others which properly think of metaphor as a linguistic phenomenon, nor does it contribute anything special to ways already canvassed in which a dead metaphor can be revived. Against this background, the fact that a group of related metaphors, say those grounded on the ways in which money qualifies time, can be brought to life together is of no special significance.

This conclusion overlaps with, and can be reinforced by, the comments made in section 4.3 about extended metaphor. What I argued there was that we can understand extensions to a metaphor as requiring multiple semantic descents to what is effectively the same object. In the example used, namely: When questioned, he offered his usual soap-bubble reason for what he had done. But this was pricked by the detective asking about the telephone call to the bank manager, and it burst completely when the bank official denied that any such call took place, an initial descent from ‘soap-bubble’ to a soap-bubble is followed by further descents from ‘pricked’ and ‘burst’ to that same item. What leads us in this case to think of a reason as something that can be ‘pricked’ and ‘burst’? Not some deep conceptual truth that reasons are soap-bubbles, but simply the fact that the author of the text led off by using a soap-bubble as a qualifier of the reason; our knowledge of the world—knowledge that bubbles are the sorts of thing that are delicate and can be pricked and burst—does the rest.

In the cases favoured by the cognitive account, we have this same kind of metaphor extension, with these two differences: the starting points are implicit, and the extensions tend to be dead or moribund metaphors. We find ourselves speaking of investing, wasting, spending time, and can obviously do so without being aware of any initial or central case of semantic descent and qualification. This can make it seem as if these apparent polysemes must be connected by some fundamental conceptual fact. But the truth is simpler. Underlying the extended metaphor that gives us these many ways of speaking is an implicit range of qualification instances: monetary transactions—these events or, in my sense, objects—can be used to qualify certain of those events in our lives that involve the passage of time. Call this fact a conceptual truth if you like, though this is just as misleading as calling it a conceptual metaphor. It is our perfectly ordinary knowledge of monetary transactions which allows us to use them for conveying information about events involving time, and it is around this basic scheme of qualification that a whole series of expressions have clustered.

4.4.5. Semantic descent and dead metaphor

My account posits two different processes in metaphor: semantic descent to an object, and the use of that object as qualifier. That there are these two is in contrast to Content Sufficient accounts in which the generation of metaphorical meanings is the hoped for result of either an alchemical combination of commonplaces associated with predicates, or an open-ended exploration of properties putatively shared by the relevant items in a metaphor. While acknowledging that neither descent nor qualification are all that tightly constrained, their remit is nonetheless more determinate than that of ‘alchemy’ and ‘exploration’. In this section, I want to consider further consequences of these two different processes for the phenomenon of dead metaphor.

As we have seen, in a dead metaphor, a meaning which is not the literal one is nonetheless available to speakers of the language, with no special effort on their

part, that is, with no apparent need for a further theoretical story on our part. The intelligibility of Romeo's assertion: (7) Capulet is a fox, does not depend on anything more than would be required of: (9) Capulet is a man. Still, dead though it is, (7) there is evidence that it is a metaphor because, unlike (9), contexts are imaginable which would force one to do more work on that same predicate than is needed for (7). Revivified in the right context, 'fox' might call on all the resources of an account of metaphor; in terms of my account this would mean calling on semantic descent and qualification to make intelligible the imagined live 'fox' metaphor.

All that is obvious enough, but, as has been pointed out by many writers, even without its predicate being brought to life in some richer context, there is something special about (7): the second, non-literal meaning may figure in a dictionary, but the linguistic act of uttering it does not leave the fox completely behind. The contrast with (9) makes the point succinctly: in accounting for the act made by an assertion of this sentence, there is no need to go beyond what the dictionary tells us of 'man', and we do not expect a hearer to treat this as other than a particular sound with that meaning. But, though the dictionary has an entry for 'fox' that includes 'crafty person', we would nonetheless think it odd to count this word as mastered by someone who simply treated the sound 'fox' as standing for a craftiness; someone who hadn't a clue about foxes as small reddish furry mammals. And yet, in trying to explain this, we certainly do not want to go back on the admission that (7) is a dead metaphor, one that doesn't call on the full range of processes of interpretation.

Would it be enough to think that the difference between (7) and (9), and the fact that (9) is a dead *metaphor*, depends simply on the counterfactual possibility described above? That is, might we simply claim that the difference lies in the fact that, in linguistic circumstances other than the straight assertion of (7), its predicate could be revivified, whereas no such thing even makes sense in respect of (9)? Obviously enough, the claim implicit in these questions is true, but it is not wholly explanatory. Even in the case where no context makes the predicate in (7) come alive, there is something of the fox in it that cannot be ignored, and any claim about metaphorical status should have a way of taking this into account.¹⁴⁰

It is at this point that the duality in my account makes a real contribution. Think of an ordering (not so much a scale) arranged as follows:

- (i) expressions (as used) whose understanding does not encourage either semantic descent or qualification, and for which these processes seem not even to make sense;

¹⁴⁰ Another way to put the point is this: the counterfactual possibility of revivification itself requires some explanation. Like counterfactuals generally, there is always a question about their categorical basis.

- (ii) expressions (as used) which do not encourage descent or qualification, but for which these processes at least make sense;
- (iii) expressions (as used) which encourage semantic descent, and allow a kind of qualification;
- (iv) expressions (as used) which encourage both semantic descent and qualification.

The range of examples I have in mind for (i) is typified by words like ‘consider’, words which might have some etymological claim to metaphoricality, but which, in current usage, neither encourage the processes of metaphor, nor seem to give them room to operate. ‘Consider’ in its original Latin home might once have suggested ‘looking at the heavens’, and it is not difficult to imagine an early Roman finding this observational pose a perfect qualifier of someone thinking over a problem. But that is not how it is with the English word. Indeed, ‘consider’ seems to be a good candidate for being dead in the absolute sense, something that is metaphor no longer.

A example for second category would be the word ‘reflect’. I would guess that this word is used without the need of semantic descent and qualification, but it is easy enough to imagine circumstances in which one might appeal to the physico-optical sense of ‘reflect’, while speaking about cognition. Still, even though such an appeal is possible, this is too slim a basis on which to declare ‘reflect’ merely a tired metaphor. If used in the physico-optical sense, but intended to bear on cognition, I imagine that the effect would be one of emphasis, the making of a forceful point, rather than metaphor. So, if the examples in (i) count as ex-metaphors, then, even if less definitively, those in (ii) do so as well.

Categories (i) and (ii) are useful in that they help sort out otherwise troubling examples, but the real novelty comes with category (iii). It is the possibility outlined there that gives us a way to have our cake and eat it in respect of ‘fox’.¹⁴¹ To see how this could work, begin by recalling what was said in Chapter 3 about the way my account handles transparency. In that discussion, I speculated that, when we come across a word like ‘sun’ or ‘fox’, it might be true that we access the linguistic meanings necessary to understand sentences in which these words occur at the same time as we access the object (a determinate object) to which these words apply. The psychological model to think about here is the one which credits us with immediate access to both of the linguistic meanings of an ambiguous word like ‘bug’, leaving disambiguation for later on in the process of comprehension. Of course, as I noted, there are several reasons to regard this as only a model: ‘bug’ has two *linguistic*

¹⁴¹ I deliberately employed two examples from Cooper 1986. In his discussion, appeal to etymology suggests just how difficult it is to distinguish ‘consider’ from a case of ‘real’ metaphor, and he introduces ‘reflect’ to show that we cannot avoid this consequence by declaring that the original meaning of ‘consider’ has simply disappeared. After all, the physical-optical meaning of ‘reflect’ is still around. My ordering shows how these two are different, without committing us to thinking that ‘reflect’ is any more metaphorical than ‘consider’. In the case of ‘consider’, there is simply no route back to the literal meaning, and in the case of ‘reflect’, we can be made to recognize the route, but no appeal to it is necessary to make our utterances with this word intelligible, even were someone to highlight ‘reflect’ in its cognitive usage by adverting to the physico-optical sense.

meanings, ‘sun’ doesn’t; the research on the processing of ambiguity concerns only the former kind of case; and finally I do not intend the semantic descent account as an armchair theory of psycholinguistics. Still, there is no harm in using something empirical merely as a model for a philosophical account, and that is what I intend. (If it were the case empirically that it also worked the other way around, so much the better.)

Here then is my story about ‘fox’. In any use of this word, two things are made available: the fact of membership of the natural kind *fox*, that is, the linguistic meaning of the predicate expression ‘fox’, and an exemplar of the animal itself. The second of these is of course the result of semantic descent. If the word ‘fox’ in (7) was used in another setting, one which revived it, the exemplar would be the focus of full-scale qualification: we would understand this word in the new setting as requiring a hearer to treat the fox as a proto-predicate of the subject of the sentence. But, as I have suggested, relying wholly on what is only the counterfactual possibility of revivification is not enough fully to justify the metaphorical status of ‘fox’ in the plain (7). Not having *actually* to call on qualification to make (7) intelligible, why should we regard (7) as having any claim to metaphoricality? The second part of my story about ‘fox’ comes in here. In cases where the exemplar doesn’t lead to full-scale proto-predication, it still has a job of work to do: it links to the dictionary meaning of ‘fox’; it functions as a sort of non-linguistic cross-referential device, and in so doing it actually short-circuits full-scale qualification. These points are apt to be puzzling for several reasons, so more needs to be said.

Anticipating (briefly) the topic of section 4.6.3, and at the risk of adding to your puzzlement, let me first ask you to think about metonymy. A standard example, has someone saying: ‘The White House decided ...’, and our realizing (these days, perhaps, only believing) that it is the US President who is being spoken about. A perfectly standard way to describe what is happening is that a certain building (The White House) is being used to refer to a person. How does this happen? I shall discuss this at some length in the later section, but it shouldn’t strain credulity to accept that there is a relationship between The White House and its principal occupant which makes such reference possible. Note too that this kind of metonymic reference is often thought to be, or at least thought to be linked to, metaphor. What I am suggesting in respect of ‘fox’ is that the exemplar, which is, as it were, present along with the word, effects reference to the second, non-literal meaning that is often given in dictionaries. Or, since the idea of a meaning as something that can be referred to in this way is not perspicuous, we perhaps should think that the exemplar refers to the property that the dictionary entry describes. There is thus a kind of cross-reference between one thing (an exemplar of ‘fox’) and another (the property of being crafty). Because of this cross-reference there is no need to appeal to the idea that the fox itself directly qualifies Capulet in (7). But then again because this kind of reference has a claim to be included under the heading ‘metaphor’—details of this in section 4.6—this would explain in part why it seems reasonable to think that a dead metaphor is still a metaphor.

What makes this explanation particularly satisfying is that it shows plainly why we are so tempted to regard even the bare possibility of revivification as a ground

for treating dead metaphors as metaphors. Metonymy works through whatever salient relationship there is between object and referent, and this is true also of the ‘fox’ case. The ground of the relationship in the latter case is this: an exemplar of ‘fox’ (a fox) can be intelligibly used to qualify human beings. Qualification in the full-scale sense is not necessary for understanding (7), but it nonetheless exerts leverage over the relationship between a fox-exemplar and the property of a human being to which that exemplar metonymically refers. What we have then is a neat story about why the fox is present even in completely ‘dead’ applications of ‘fox’ to human beings, a story that gives the fox a plausibly metaphorical role as a metonym, while still managing to include the full-scale qualification that, in richer contexts, is actually needed.

Note too that this way of treating dead metaphors avoids the problem of the missing fox that plagues other Content Sufficient accounts: the fox—or an exemplar of that kind—is the crucial link between the word ‘fox’ and the established second meaning of this word. If it were to happen that we came to use some other word for the natural kind *fox*, and that the second meaning of our ‘fox’ came to be its only meaning, then ‘fox’ would move up the ordering to category (ii) or even possibly (i). There would no longer be any metonymic cross-reference. We would in this case have a metaphor that, like the parrot, was genuinely ‘ex’—an expression for which there would be *no way back* to the processes of metaphor.

My story differs crucially from any told by other Content Sufficient theorists. Their starting point, like mine, is the *word* ‘fox’. Their explanation of the predicate in (7) being dead is that the word ‘fox’ has come to have a second meaning. In effect, ‘fox’ has become ambiguous, it resembles ‘bank’. Of course, it only *sort of* resembles this straightforwardly ambiguous word. So, we now have to explain how the ambiguity of ‘fox’ differs from ‘bank’. Mostly commonly this is done by adverting to differences in the histories of these two words: ‘fox’ somehow ended up with its two meanings because of some previous metaphorical process, whereas ‘bank’ never had metaphor in its past. Note that only genuinely historical metaphorical processes are allowed here. For, if we allowed that contemporary English speakers engaged in them, then ‘fox’ in (7) would be a live, not a dead, metaphor.

On my account, there are two meanings at issue, but one is tied to the word ‘fox’ and the other to the fox itself. The exemplar that hovers around the ‘fox’—the item of semantic descent—is what links the word to the second meaning. There is no equivalent to this in cases like ‘bank’, and we thus have a principled way to distinguish dead metaphors from ordinarily ambiguous words. Moreover, unlike the usual Content Sufficient story, we do not have to appeal to some imagined *history* of ‘fox’ to make the live/dead distinction. We can explain the role of the fox in our understanding of (7) by appeal to semantic descent, we can explain the deadness of (7) by the fact that the fox exemplar is not called upon to qualify Capulet, and we can explain its metaphoricality by the fact that the fox metonymically picks out the second meaning, and does so because of the present, not historical, possibility of using a fox qualificationally.

4.4.6. Idioms, live metaphors, and dead predicates

There are further consequences of the ordering (i)–(iv). In this subsection, I discuss several of them. Along the way, I should also like to clear up what may be some residual confusion about the conditions required for (a use of) some expression to be a metaphor.

(a) I have said nothing detailed about category (iv). In one way, this is not surprising: items in this category are live metaphors, so (iv) really marks the other side of a boundary, the limit beyond which there is nothing deserving the label ‘dead metaphor’. So, the details of (iv) are simply not relevant to a discussion of dead metaphor. Still, I have included it in the ordering, not simply as a boundary, but with the idea of preparing the ground for the discussion of richness in metaphor that I am working towards. Let me explain.

The crucial point about the categories extending from (i) to (iii) is the fact that they differ in a *principled* way from one another; they are not simply nodes on a scale along some dimension like currency or vivacity. Category (i) consists of expressions for which semantic descent and qualification are neither needed nor appropriate. I used ‘consider’ as an example here, but the category forms a catchment for a number of otherwise problematic cases. Thus, idioms, for example, ‘kicked the bucket’ or ‘by and large’, are naturally at home here.¹⁴² In the case of expressions like ‘kicked the bucket’, the striking descriptive element may well make one wonder how they came to be established, but this research does not really add anything to their meaning. Any investigation *begins* with the fact that the expression means ‘died’, and the only thing at issue is how it came about that ‘kicked the bucket’ means that.¹⁴³ In showing how category (i) deals with idioms, the second of my examples is perhaps the better one. English speakers are by and large not even curious about why ‘by and large’ means ‘generally’, so the fact that it is originally a nautical expression can come as a surprise. Sailing ‘by’ the wind is sailing with the wind ahead, and ‘at large’ is sailing with the wind behind. Hence, ‘by and large’ covers pretty much every direction of sailing and wind. Would learning this origin be more than a curiosity? Once learnt, do we think that ‘by and large’ could be revived and is therefore a dead metaphor? I think not. For the crew of a seventeenth-century sailing ship, maybe there was something of the metaphor in the expression, perhaps even a live one. But even knowing this about its origins is not enough to allow us to demand either or both of semantic descent and qualification. Put in less psychological terms, there is neither need nor place for descent and qualification in any account we give of the intelligibility of *any* sentence that uses ‘by and large’.

This is in contrast to expressions in category (ii) where the processes of metaphor make sense, even though they are not needed to understand relevant expressions.

¹⁴² There is, however, no point in being inflexible here. It may well be that a case-by-case examination would find some idioms more suited to category (ii). If so, fine, even interesting, but it will still be the case that idioms are not metaphors, and that is the crucial point made by my classification.

¹⁴³ The grammatical inflexibility of ‘kicked the bucket’ is another among the signs that we are not dealing with a dead metaphor; one cannot speak of someone on their death-bed as kicking the bucket. Hence, for all that it is a picturesque expression, revivification is simply not on the cards.

'Reflect' goes about its business in various contexts without our being called upon to connect its cognitive and optical senses. We can of course devise sentences in which the connection is made obvious, and there is no doubt that the two senses are related, but the connection itself isn't enough for the possibility of revivification. Thus: The thought bounced back and forth between them, reflected but not actually reflected upon, links the two senses, but that is all. It is not that we learn anything about the one sense from the other, nor are we led from one to other as we are in cases like 'fox' or 'mouth'.

Examples like 'reflect' are only the tip of the iceberg as far as category (ii) is concerned: think about the vast range of prepositions which seem to possess concrete and, as is all too quickly said, metaphorical senses. An airplane may be said to be *on* time, and this may well have some connection with our also saying that the cup is *on* the table, but the use of 'on' in both senses, as in: The airplane was on time and on the runway, does not breathe any kind of life into what some suppose to be a metaphor here. Familiarly, the simultaneous use of different senses of the same preposition, as in this from Lewis Carroll's *The Hunting of the Snark*: They sought it with thimbles, they sought it with care, count as instances of zeugma, one of the traditional tropes, but no one thinks that zeugma is itself a species of metaphor.

With categories (iii) and (iv) we come to genuine instances of metaphor, first dead and then alive. I won't repeat what was said above about these, but I would like to put a marker down here that will be important later. If we could devise a principled way to sharpen the ordering in the vast and unruly category (iv), then we would have a better chance of understanding the kind of thing that writers mean when they speak of metaphors as more or less vivid, 'high-octane', or rich. And given that what separates category (iv) from category (iii) is that the former actually calls on objects to serve as full-scale qualifiers, the obvious way to extend the principled ordering is by making finer grained distinctions within this activity. For example, one might distinguish culturally determined uses of objects of qualification—the sun as used of Juliet might be an example here—from those which depend more on immediate contextual saliences, for example, Kant's claim about crooked timber. Without wanting to anticipate too much of what will be said in the final section of this chapter, it might well be thought that any kind of established use of an object as a qualifier is likely to go with a less vivid, though of course still live, metaphor.

(b) The second consequence of the ordering will actually figure in the next section. Still, since it is dependent on my remarks about dead metaphors, and since it might

be surprising, it seemed a good idea to prepare the ground in advance of my use of it in the discussion of simile and metaphor.

Words in category (ii) like 'reflect' might once have been metaphors (in some language and culture important in the development of ours), and they still possess multiple senses. But they are not now even dead metaphors: they are beyond revivification, and their multiple meanings do not have the requisite linkages. What I should like to suggest here is that there is a case for putting ordinary predicate expressions like 'ewe' into category (ii). This is the surprising contention referred to in the previous paragraph, and I hope that what follows will at least remove some of the initial counter-intuitiveness that is no doubt responsible for the element of surprise.

It is important not to be misled straightaway into taking my claim to be that ordinary predicates once were, or might have been, metaphors. This is certainly true of 'reflect', but most certainly not true of 'ewe'. Still, I do maintain that the rubric of category (ii): 'expressions (as used) which do not encourage descent or qualification, but for which these processes at least make sense', can be understood in such a way as to make 'ewes' welcome. And that they belong in this way to (ii) will be an ingredient in my proposal about the relationship between simile and metaphor. However, putting my firm insistence that ordinary predicates are not metaphors alongside my many scattered remarks about what counts as a metaphor, dead or alive, I think I should begin by pulling together previous explicit as well as implicit remarks on this general topic.

One ultimate aim of my account is to make it plausible that metaphor constitutes a kind; that there is some underlying feature which will allow us, not perhaps without some vagueness, to group together as metaphors all and only those sentences, written or uttered, with the relevant feature. I say 'ultimate' here because I don't take myself to have done more than lay the groundwork for this. What I have claimed is that there are two functions which words in various sentences (as used) might take on, and which can help us understand why these sentences are intelligible. These are semantic descent and qualification. My view is that when we need to appeal to both of these functions in order to understand the intelligibility of some sentence or larger unit of discourse, then this is a sure sign that we are dealing with metaphor.

This is all too easily misunderstood. I am not claiming that speakers and hearers have recourse to these functions, and that their doing so is a *symptom* of metaphor. I have throughout been completely agnostic about how metaphors are identified, and for all that my exposition might at points encourage it, I do not intend these functions to be present in the minds of speakers and hearers. Rather, the claim is that, when and in whatever way we do identify a metaphor, it will be by reference to semantic descent and qualification that can render the linguistic act of uttering or writing the relevant sentence or sentences intelligible. Putting it this way should be enough to forestall various misunderstandings.

One should not think of the examples in Chapter 2 as metaphors. What I sought to provide in that chapter was some reason to think that we should liberate the

notion of predication, so that it was not a function restricted to devices in natural language. Having done this, I then offered some evidence by example that there really are cases which fit this newly freed notion. Still, cases like Jones's encounter with the wind-felled tree, and Nabokov's with the bizarre paling, were not themselves metaphors. As I asked you to think about them, these objects were not brought to the attention of the relevant witnesses by linguistic means; though they were examples of qualification, they did not call for anything like semantic descent to make them intelligible. There were, as I imagined it, no words or sentences uttered for which the question of intelligibility even arose. Of course, in so far as you do see qualification at work in these examples, it is a short step to imagining them offered linguistically to convey information to others or to the subjects themselves. Indeed, it takes some effort to resist doing this, especially in the Nabokov case, where the example figured in his memoirs. When there are sentences used in this way then I think we do have cases of metaphor, cases in which the utterance of some sentence is only intelligible by appeal to a story about metaphor, and, in particular, my story about descent and qualification.

Is it then my view that whenever we have recourse to semantic descent and qualification we have metaphor? The ordering presented in the previous subsection might suggest this, but care is needed here. There can be cases in which semantic descent and qualification could fit a sentence, without this making the sentence or any of its constituents metaphors. One such case was the sentence: Einstein is the brilliant scientist. I didn't expand on the point in Chapter 3, but I did note that, for all that we could speak of descent and qualification here, even understanding the sentence in the way proposed did not make it a metaphor. The reason for this should now be clearer. It may well be that this sentence can make us think of an exemplar of a brilliant scientist, and that we then use that object as a qualifier of Einstein. I encouraged this way of looking at the sentence as part of dealing with the surprisingly intricate structure of Romeo's apparently straightforward (R). But we do not have to have recourse to descent and qualification to make the Einstein sentence intelligible; the sentence is perfectly intelligible as it is, and differs in this respect from (R). Admittedly, we could take the predicate expression in the Einstein sentence to link us to certain properties—those of brilliant scientists—rather as we did with 'fox'. But this linkage, even resembling as it does the metonymic linkage between 'fox' and the property of craftiness, is not enough to force the verdict that there is something metaphorical about it. What is in question in the Einstein sentence are properties of certain human beings, and both Einstein and an exemplar of 'brilliant scientist' can possess those properties without this being even slightly puzzling. Whereas some explanation of the fact that English speakers use foxes in the way they do is called for, even in this most dead metaphor.

What I am saying, then, is this: there is a category of sentences or constituents of sentences (as used on specific occasions) which cannot be made intelligible by

the usual meaning-theoretic tools at our disposal. But they can be made intelligible by recourse to semantic descent and qualification. What I think follows from this—and I will offer further support in succeeding pages—is that this category is in fact the kind *metaphor*. Metaphor is thus an explanatory kind: its constituents are unified by their requiring and getting a reasonably well-unified kind of explanation of their intelligibility.¹⁴⁴

Against this background, it might now be clearer why I described category (ii) words like ‘reflect’ in the way I did. These are not metaphors because their use is perfectly intelligible without recourse to semantic descent and qualification, but they are different from ‘consider’ and ‘by and large’, in that descent and qualification make sense in relation to them. What I should like to describe now is why this is also true, even of ordinary predicate expressions like ‘ewe’.

The place to begin is with the just-so story that I told about the origin of predicate expressions, though there will be no need to rehearse all the details of that story here. What matter in the present context are only its starting and, most particularly, its end points. The starting point was the idea that qualification plays a crucial role in the origin of familiar categories. It is because one ewe is informative about, can qualify, another, and because the population of ewes happens to allow mutual qualification, that a practice develops in which one can imagine, first, there being a Standard Ewe, then, a concrete token adverting to this Standard, and, finally, a sound (i.e. word) replacing the token. While I allowed it to be true that ewes share a property or properties, this fact was not found to be explanatory. What I tried to get you to see was that this sharing of properties is consequent on some more primitive practice, indeed a practice that begins with qualificational insights, more primitive, but not essentially different from those of Jones and Nabokov.

The use of a sound, say ‘ewe’, as a replacement for a token, itself a replacement for the Standard, was not quite the end point of my story. For I do not really think that the sound ‘ewe’ serves now primarily as a replacement for these things, but rather as a replacement for the information which grounds the possibility of mutual qualification. ‘Ewe’, in one of its uses, picks out whatever it is which makes it true that these creatures are mutually qualifying, a fact we can perspicuously describe by saying that ‘ewe’ picks out the property of *being a ewe*. In this way, one can see both how we come to talk of properties, and also how easy it is to overlook the role that qualification plays in the development of that talk. My point in Chapter 2 was that in so far as there is something in the just-so story, one can see why qualification can seem a rather rarefied phenomenon: ubiquitous though it is in getting us started on the road to categorization, it drops out of sight once these categories are established.

A further point, itself some slight evidence for my story about categorization, is that while it has dropped out of sight, qualification is not wholly missing, even from our current uses of ‘ewe’. My idea was that we can see it at work, behind the scenes, in the tendency of philosophers (and others) to use ‘ewe’ and ‘ewehood’ interchangeably, even though, on reflection, everyone can see that it is illicit to speak of ‘the property

¹⁴⁴ I shall suggest later that some of the ways we actually use ‘metaphor’ give evidence of this. While I do not insist that my proposal wholly captures the intuitive or everyday use of ‘metaphor’, I do not mean it to be completely revisionary.

of ewe' or to use 'ewehood' in places where 'ewe' is called for. Without repeating what I said there, an ingredient in my explanation of this was that we could imagine applying semantic descent even to the most ordinary of predications. To remind you of how this works, consider: Clio is a ↓ewe↓. Though nothing about it demands this treatment, we can, perhaps surprisingly, make sense of the semantic descent marked here, but only if we also avail ourselves of qualification. In effect, the sentence is tantamount to: Clio is a *this* (*here the speaker demonstrates an exemplar of 'ewe'*). And, as we saw in Chapter 3, this makes sense only if we take the object demonstrated to supply information about Clio, as in a predication—to be in fact a case of qualification.

It is thus a consequence of my story that, even though we have to look hard for it, there is evidence that the processes of metaphor make a kind of sense, even in cases of perfectly ordinary literal predication. It is this conclusion that lies behind my including ordinary predicates in category (ii).

As noted earlier, I can imagine surprise at the inclusion of perfectly ordinary literal expressions like 'ewe' in the same category as words like 'reflect'. But surprise here should not be allowed to turn into puzzlement. I am *not* saying that 'ewe' is an ex-metaphor. Indeed, strictly speaking 'ewe' was never a metaphor, because at the stage when categorization got going, only qualification figured. The use of 'ewe' (in the just-so story) *followed* various qualificational insights and the establishment of a practice, so it simply couldn't have figured in a process of semantic descent. In this way, there is almost certainly a difference between 'ewe' and 'reflect'. The possibility of saying that we can 'reflect' on some subject matter may well have come about by properly metaphorical processes: early Latin speakers may have pressed a physico-optical exemplar of reflection into service in characterizing a certain mental process, and we have adopted a cognate form of their word to serve these two purposes. Clearly, there is nothing like this true of 'ewe'. However, while there is this difference, there is a much more significant similarity: for all that we can apply them to sentences involving 'reflect' and 'ewe', semantic descent and qualification are not required for intelligibility. 'Ewe' never was a metaphor, and a word like 'reflect' is not even a dead one.

The fact that we can regard the development of ordinary predicates as at least sharing a process with metaphors, and that, even in current use, the processes of metaphor could be applied to them, should be more satisfying than puzzling, especially to those who think of metaphor as in some sense fundamental to language. Though I will return to this in the last section of the chapter when I consider the 'richness' of metaphor, it should be easy enough to anticipate what might be said there. What could be more indicative of the fundamental nature of metaphor than its calling on at least one of the processes that also figure in the development of ordinary predicate expressions? Traditionally, many writers have suggested that metaphor

does figure in the origin of many ordinary, non-metaphorical words, but there has been a tendency to support this with an appeal to ex-metaphors, idioms, and dead metaphors. We are expected to be impressed by the fact that expressions in everyday literal use were once metaphorical. However, the appeal to these is often too indiscriminate and superficial to carry much conviction. It all too readily encourages a response something like this: sure, there are expressions here and there where metaphor seems to play a role. But many of these expressions are not actually metaphorical, and those that have some claim to this label scarcely suffice to make a case for the fundamental nature of metaphor in the development of language.¹⁴⁵ In contrast, I have in this section suggested that we can sort expressions into those which have had some kind of claim to metaphoricality, and those which no longer do, precisely by focusing on the use of the processes that by my lights are required by metaphor. These processes are not necessary to explain the linguistic acts effected with 'ewe', 'reflect', or 'by and large', so these acts most definitely do not count as metaphors, but the fact that they make the kind of sense they do in regard to category (ii) expressions highlights the importance of metaphor in a quite different way.

4.5. Simile and Metaphor

Perhaps it will be obvious why discussions of mixed, extended, and dead metaphor are relevant to my promised characterization of richness in metaphor. But while the subject matter of this section, and the two which follow, do not so obviously tend in this direction, they will in fact be material to that characterization. Also, dealing as they do with certain central phenomena of metaphor, they have an importance independent of the longer term goal of this chapter.

4.5.1. The simile view

A certain view about the relationship between metaphor and simile is perennial. In the first, spring-like rush of optimism, it can seem obviously right to maintain straight out that metaphor is an elliptical or compressed simile. We save ourselves lots of trouble, for example, by treating Romeo's remark as equivalent to its 'expanded' form: (Rs) Juliet is like the sun. After all, (Rs) seems to convey pretty much the same message as our original (R), it could well be true—certainly it is a candidate for truth—and it offers the prospect of handling features such as, among others, transparency.

¹⁴⁵ This is a good place to insert a further remark about the cognitive (or conceptual) account of metaphor. Lakoff and Johnson, and those who think well of their views, feel strongly that the view shows how extensively metaphor figures in natural language. Others tend, just as strongly, to dismiss the view, often on the ground that a lot of what they claim as metaphorical is no such thing. Earlier in this chapter, I suggested that the key element of the cognitive account—the idea of metaphor in thought—is really a rather loosely described version of my notion of qualification. Seen in this way, we can understand why there is something right about the aspiration of the cognitive account in respect of metaphor, as well as why the cognitive account doesn't handle this aspiration correctly.

However, by autumn, the view tends to look less attractive. Many are put off by the fact that in some sense everything is like everything else, so it is difficult to see how the ‘false’ (R) can be equivalent to (Rs), and, aside from this, any information conveyed by (Rs) is unimpressively truistic. Finally, the move from the problematic (R) to the all too obvious (Rs) suggests that paraphrasing a metaphor is not only possible, it is implausibly easy.

Still, being a perennial, I have no doubt but that the view will come to seem right to some writers on metaphor. It will be said, with real justice, that many of the arguments against it are too quick, and it will again come to be thought a rival to other accounts of metaphor.¹⁴⁶

Views that, in one way or another, resurface again and again are particularly important in philosophy: we often learn more by coming to understand why they are perennial than by trying to decide whether on balance they should be accepted or rejected. This seems to me especially true of the simile account of metaphor. In what follows, I shall try to explain why the simile view refuses to go away, why it is perennial. But even perennials have a finite lifetime, and I suggest that the time has come to dig it out and replace it. The replacement offered will resemble the simile view closely enough to satisfy any demand for continuity, but being nurtured by the semantic descent account, I shall argue that it is an altogether more robust specimen.

The simple fact that accounts for the simile view's perennial appeal is that it says something true. Or, more accurately, when it is construed correctly, and when it is allowed to operate in a certain narrow sphere, a claim central to the view is true. I will begin with a general statement of that central claim, and then consider the necessary pair of qualifications needed for the plausibility of that claim. Reasons for replacing the simile view in spite of the truth it contains will come later.

When someone says of the adult Tolstoy:(T) Tolstoy is an infant,this suggests a comparison between Tolstoy and an infant, and it further suggests that there are properties shared by Tolstoy and some typical infant. As I have argued on several occasions (and I will return to this in the next chapter), the fact that comparison and property-sharing are in the air when someone asserts (T) is not enough to account for its intelligibility; it is not itself an explanation of metaphor. But, having said this, I do not deny that (T) implies comparison. And this is the kernel of truth that makes the simile view perennial. For, having found (T) to be comparative, the focus shifts to:(Ts) Tolstoy is like an infant.The word ‘like’ is almost universally taken to be a marker of comparison, so, given that (T) suggests comparison, it is certainly tempting to think that (Ts) makes explicit what was only implicit in (T). Succumbing to this temptation is, in effect, accepting the simile view—the view that (T) is an elliptical way of saying (Ts).

¹⁴⁶ Fogelin 1988 devoted himself to righting the injustices done to the simile view, and it is his defence of this view that will figure later in my discussion.

Now the first of the problems arises right at this point. In judging ‘like’ to be a marker of comparisons, as in (Ts), we must not forget comparisons of the common-or-garden sort. Thus, in telling someone who has done very little carpentry: (11) A screw is like a nail, we are apparently comparing the two kinds of fastener. But we would not regard (11) as a non-elliptical version of: (11m) A screw is a nail, since the latter is not only flatly false, it is hopeless as a way of capturing the straightforward thought expressed by (11).

What is pretty much the same problem viewed from a different angle arises with metaphors like (T). We may want to treat the non-elliptical (Ts) as equivalent to (T), but if we have the screw-and-nail comparison in mind we won't have any such right. The claim in (11) leads us to expect there to be some central feature actually possessed by a screw that is no less actually possessed by a nail. But we cannot treat (Ts) as a straightforward comparison on the model of (11). Tolstoy and the infant share many perfectly ordinary properties: they are both material objects, both human beings, both have eyes, ears, etc. But it would be absurd to regard any of these as explaining the supposedly elliptical (T).

The issues raised in the preceding paragraphs are two sides of the same coin: they arise from our regarding *any* comparison made using ‘like’ as an explicit or unelliptical version of an implicit or elliptical comparison made without ‘like’. What causes the trouble in the case of (11) is that there is something strange about the implicit comparison, and what causes trouble in the case of (T) is that there is something strange about the explicit comparison.

Obviously, the thing to do is to make a distinction between *kinds* of comparison. There are two ways to do this. On the one hand, we could insist that ‘like’ is not strictly univocal: it is sometimes a marker of literal comparison and sometimes of non-literal.¹⁴⁷ On the other hand, we could treat ‘like’ as univocal, and treat the distinction between kinds of comparison as a matter of context: ‘like’ always marks comparison, but pragmatic factors indicate whether the comparisons so marked are literal or non-literal. In whatever way we manage it, some such distinction would give us room to manoeuvre. Using this room we can say that, when it is a matter of literal comparison, no elliptical form is possible, and that when it is a matter of non-literal comparison, as in a metaphor, the non-elliptical form must not be misinterpreted as a literal comparison. In terms of our examples:

- (i) to say that a screw is like a nail is to compare two things *literally*—perhaps it is to use the literal ‘like’—and there is no counterpart elliptical version of this;

¹⁴⁷ Regarding ‘metaphor’ as the name of an explanatory kind, I should be more careful about lumping it together with the merely ‘non-literal’ or ‘figurative’. However, my exposition of Fogelin will be easier to follow, and no permanent damage done, by this carelessness. See section 4.7 for more on ‘figurative’.

- (ii) to say Tolstoy is an infant is to compare them *non-literally*, and we must therefore think of the explicit version—the one employing ‘like’—as also non-literal.¹⁴⁸

Of course, though making a distinction between two types of comparison allows us to say these things, the only justification for the distinction so far given is that it gets the simile account out of trouble. In any full treatment, we would have to be told why there happens to be no appropriate elliptical version of a literal comparison, and what is involved in comparing two things non-literally. But, since my aim is ultimately to replace the simile view, there is no point here in considering detailed justifications for the simile view's attachment to kinds of comparison. However, because it figures later on, I do need to say something about a sort of pecking order that exists amongst these comparisons.

It is central to the simile view that literal comparison statements like (11) are in some sense basic, and that non-literal or figurative comparisons are parasitic on them. Fogelin writes: ‘Figurative meaning arises, *in general*, through a (mutually recognized) mismatch of literal meaning with context, and, more specifically, this is how the figurativeness of figurative comparisons arises’ (Fogelin 1988: 30). And we can take him as saying that, on encountering a claim of this form: (12) A is like B, and finding that it makes little or no sense, when taken in the ‘ordinary’ or literal way, we should reconstrue it as, say: (13) A is figuratively-like B. Fogelin's summary of the simile view follows straightaway: ‘The figurative meaning of a metaphor of the form “A is Θ is the same as the *figurative* meaning of the counterpart simile of the form “A is like Θ (Fogelin 1988: 29).¹⁴⁹

It is just at this point that the second problem with the simile view emerges. As summarized above, the simile view seems restricted to metaphors of the subject-predicate form. Fogelin is quick to note this, and he appends the following to the summary above: ‘This specification is incomplete since ... metaphors come in different forms and transpose into similes in different ways, but these simple patterns

¹⁴⁸ The cornerstone of Fogelin's defence of the simile view is his insistence on the distinction between literal and figurative comparisons. Anyone who has read that work will see how it has shaped my discussion, but, except for a few passages I am about to consider, I shall reserve comments about Fogelin's overall view until the next chapter.

¹⁴⁹ It is part of Fogelin's view that the literal meaning of (T) is the same as the literal meaning of (Ts). It is certainly true that both (T) and (Ts) are bizarre when we try to understand them literally, but it is far from clear that they are bizarre in the same way. In any case, the equivalence of literal forms is not required by Fogelin's version of the simile view itself, but is rather a consequence of a general claim about ellipsis that he defends, namely the thesis that if any expression is elliptical for another, the two must have the same literal meaning. I don't know if this thesis is defensible, but it won't be necessary to consider it here. Still, when I come to give my ‘replacement’ for the simile view, one surprising consequence will be that a version of this equivalence thesis will come out true.

will serve my present purpose.’ However, this is an (implicit) promissory note that is never really fulfilled so, at the very least, it would be safest to proceed as if the simile view were restricted to metaphors which are either of the subject-predicate form, or can be gerrymandered into that form.¹⁵⁰

Taking stock then of where we are:

- (a) The simile view rightly maintains that metaphors of the form ‘A is B’ and the simile form ‘A is like B’ both imply comparisons between A and B. But, though not argued at length in this section, this truth is not enough to make the simile view an explanatory account of metaphor.
- (b) Each comparison in (a) must be understood non-literally. In particular, we must not think that a metaphor is equivalent to a literal statement of the form ‘A is like B’.
- (c) The notion of literal comparison is basic: a figurative comparison, whatever it is, arises as a special case when we cannot make proper sense of a comparison as literal.
- (d) The simile view is only appropriate for subject-predicate metaphors, or for those which can be transposed into subject-predicate form without distortion.

These points explain the appeal of the simile view, especially given the tendency of discussions of metaphor to restrict themselves to subject-predicate examples. This is because it is all too easy to believe that the kernel of truth in (a)—the fact that metaphors often imply comparisons—explains what is going on when someone utters a metaphor. I have throughout this book insisted that comparison and property-sharing are not in this way explanatory, and I shall revisit this issue in the final chapter. But in the rest of this section, I should like to describe a way in which the semantic descent account can give us a different take on the simile view.

4.5.2. ‘Like’

The place to begin is with ‘like’. The simile view assumes that the primary function of this word is to mark comparison, but for all that this seems obvious, I think there is reason to question it.

Recall that when I considered the variety of syntactic forms that metaphors can take, it was suggested that ‘like’ can play the role of a hedge, softening what would otherwise be too jarring a juxtaposition between whatever item a metaphor is about, and that object which, in my account, does the metaphorical work. It was also argued that there are a whole variety of other expressions that play a similar, but often more pointed, hedging role. To reprise one of the earlier examples, compare: Faunia’s sculpted Yankee features are a narrow room with windows in it but no door,

¹⁵⁰ It would be too much of a diversion to consider this issue here, but recall the many examples already discussed earlier in this chapter which cannot be transposed into subject-predicate form, even though they have an underlying ‘this-is-that’ quality. The latter is of no help to the simile view, however, because this underlying feature is only revealed by the semantic descent account, and is not something present in the surface structure of relevant metaphors.

with:Faunia's sculpted Yankee features made me think of a narrow room with windows in it but no door.To say that the addition of 'made me think of' to the latter makes it clearer is an understatement; the sentence without it is barely interpretable. Try now:Faunia's sculpted Yankee features are like a narrow room with windows in it but no door.The effect of 'like' here is parallel to that of 'made me think'—it serves to prevent the jarring that comes from the immediate juxtaposition of facial features and room—though it is an all-purpose hedge in a way that 'made me think of' is not.

Of course, the fact that 'like' has this hedging function doesn't prevent its also being, indeed primarily being, a device marking comparison in the way required by the simile view. Still, the above example should give us pause. Remember that the simile view insists that figurative 'like' claims are *equivalent* to their elliptical (i.e. metaphorical) counterparts. Yet this seems scarcely credible in the case of Faunia's features.¹⁵¹

The issue of 'like' as a sort of hedge is only one of the problems with the simile account's handling of this word. But, rather than going through them one-by-one at this point, I think it best if I present my alternative account. Having done so, it will be much easier to appreciate the simile account's difficulties, both with 'like', and more generally.

The treatment of 'like' within the semantic descent account is straightforward, but radical. Instead of being primarily a marker of comparison, I take 'like' to be an explicit marker both of semantic descent and the qualification that it brings in train. Thus, we should understand:(Rs) Juliet is like the sun,as:(Rd) Juliet is ↓the sun↓,where 'like' is in effect an explicit instruction to engage in semantic descent, and the qualification which follows on from it. The consequences of this are far-reaching.

First, it provides a nice explanation of the hedging function of 'like'. In those cases where there could be a problem in directly juxtaposing a metaphor subject and qualifying object, employing an explicit instruction of what is to be done can be crucial. This was evident in the case of Faunia's features considered earlier. Moreover, as we have seen, 'like' is not the only such hedge, albeit it is an all-purpose one, and

¹⁵¹ The lack of equivalence is especially marked in cases where the 'hedging' is accomplished, not by some word or phrase but by the possibility of multiple interpretations and the shifting among them that I discussed in section 4.2. However, because these are also typically cases of metaphors not capturable in subject-predicate form, making them the centrepiece of my discussion at this point would have risked entangling the issue of 'like' with the issue of the syntactic restrictedness of the simile view.

my proposal allows us to collect together all such devices. One and all they include explicit instructions to engage in semantic descent, but some include more. Thus, we can take Roth's original sentence: Faunia's sculpted Yankee features made me think of a narrow room with windows in it but no door, to push us in the direction of: Faunia's sculpted Yankee features are ↓ a narrow room with windows in it but no door ↓, but to do so indirectly, by adding an explicit reference to the narrator's thought processes. Thus, this sentence is best captured as something like: Faunia's sculpted Yankee features are, to my mind, ↓ a narrow room with windows in it but no door ↓.

Second, note how this proposal can explain, without necessarily accepting, the virtually universal tendency to regard 'like' as a marker of similarity. For the simple fact is that when an object B can be used to characterize another object A—when B functions as part of a proto-predicate of A—it is bound to be true that B is in some sense similar to A. To see this clearly, begin by thinking of a case where what does the characterizing is a linguistic predicate, for example: Juliet is a woman. The predicative function in this sentence is fulfilled by a linguistic device—'is a woman'—the mastery of conventions for the use of which is part of competence in English. We do not need to look beyond this device in thinking about the truth or assertion conditions of the sentence; we do not need to think of an exemplar of 'woman' in understanding the sentence. Now, one of the central claims in this book is that this same predicative function can be exercised, not by words, but by objects themselves. And when it is so exercised, it is not the objects *qua* objects that matter, but their use as predicates. So that when we say: (R) Juliet is the sun, we are using the sun itself in a predicative way to convey information about Juliet. However, precisely because the object is made salient by what is, according to me, the semantic descent implicit in (R), there is the constant and natural temptation to overlook the predicative function, and focus on the object itself. We ask ourselves questions such as: What is the sun's relation to Juliet? Is it similar to her and, if so, in what way? These questions are perfectly in order, but we make a mistake if we think that in answering them we are explaining what Romeo is doing in uttering (R). For of course Juliet is not similar to the sun in any but a figurative sense, and the latter depends upon, and so does not explain, the more basic fact that the sun in (R) functions as a non-linguistic predicate—a qualifier—of Juliet.

Given the above, and as promised, my proposal shows why the simile view says something true (at least in its restricted subject-predicate sphere of action), even though this truth is not a full vindication of the view. For when we say: (Rs) Juliet is like the sun, and take 'like' as an explicit marker of semantic descent and qualification, we end up in roughly the same place as the simile view. (R) is indeed equivalent to (Rs), give or take the hedging that accompanies the explicit instruction in (Rs). But the equivalence here is founded on a parity of *predicative* function in both (R) and (Rs), and not on the unhelpful idea that both of these sentences claim a similarity between Juliet and the sun. The latter thought is unhelpful in two senses.

To start with, there is something hopelessly weak about seeing the whole point of (R)—and metaphors generally—as claiming nothing more than similarity. Similarity judgements are not useless for conveying information: it is often helpful to be told, for example, that one person resembles another, either physically or psychologically. But in those cases in which it is helpful, this is invariably because it calls on some contextually relevant characterization of the point of the similarity, some characterization that is explicitly given, or is presumed by the participants in the conversation. In a typical case, one must know something about B, perhaps their height, that is relevant to the similarity in order for it to be informative to be told that A is similar to B. When this prior characterization is absent, the judgement is, if intelligible, scarcely informative, and yet this is just what we are offered on the simile view. For given the things we know about the sun, it is simply false that Juliet resembles it in any way that would make the similarity judgement informative.

The second ground for finding similarity unhelpful follows straight on from this. As we have seen, since Juliet is not in fact similar to the sun, the judgement that she is must be given a special status as 'figurative'. But then the burden of explaining (R) and (Rs) is simply shifted onto the story that must be told about this status. No such burden arises if we treat metaphors, and their near-equivalent similes, as involving the special kind of object-predication that I have called qualification. Moreover, in understanding an instance of qualification in a metaphor—in being attuned to a particular use of an object as predicate—we will have laid the groundwork for any presumed figurative similarity. If an object B got by semantic descent (i.e. in a metaphor) conveys information about an object A, then, in so far as we reflect on that information, we will put ourselves in a position to see how it is that A is figuratively similar to B.

At this point, I can imagine someone puzzled about how my proposal deals with 'like' in perfectly ordinary (i.e. literal) judgements. Surely, it will be said, in such judgements, 'like' has everything to do with similarity and nothing to do with semantic descent and qualification. My reply to this constitutes the third of the radical (though beneficial) consequences of my account—one sufficiently involved to require its own subsection.

4.5.3. Literal and figurative

Consider the following claim: (12) Rugby is like football. It is difficult to avoid thinking of this as other than a straightforward comparison, and, hence, that 'like' here means 'is similar to'. However, on my proposal for 'like' in similes, this claim would be understood as: (12') Rugby is ↓football↓, where the game of football is called upon to fulfil the predicative function of giving information about, or characterizing, rugby. Now it must be admitted that my construal apparently runs counter to the rather simpler and intuitive idea that, in literal claims, 'like' just means 'is similar to'. It might therefore be best to restrict my semantic descent story about 'like' to similes. After all, as has been noted already, any plausible version of the simile view has itself to posit a special kind of 'like' judgement—the figurative—so my proposal would be no worse off. However, this doesn't seem right, and if we can avoid duplicity, so much the better. A brief reprise of the problems caused for the simile account by this kind of duplicity will make clear how I intend to proceed.

In addition to the need to posit a figurative sense of comparison, adherents of the simile view also insist that this figurative sense takes off from the more basic notion of literal comparison. Very roughly, a literal comparison such as (12) above claims that there is some salient property or properties which football shares with rugby, and of course this is true.¹⁵² In a simile, there will be no such straightforward property-sharing, and our attempt to make sense of the claim that there is results in a different kind of comparison, a figurative one. Linking the figurative sense to the literal—insisting that one develops from the other—is obviously a good defence against the charge that figurative comparison is merely ad hoc. But unfortunately the simile view encounters difficulties with the linkage, difficulties which threaten to undermine this defence.

As noted earlier, literal 'like'-claims do not have equivalent non-'like' forms: it makes no sense to say: (12'') Rugby is football. But figurative comparisons not only have non-'like' forms, the latter are the focus of the exercise. After all, these forms are metaphors, and the whole point of the simile view is to claim an equivalence between relevant 'like' constructions and metaphors. So, the unintelligibility of (12'') makes one wonder whether the figurative sense of 'like' is really based on the literal one.¹⁵³

¹⁵² This is only rough because some, e.g. Fogelin, are more careful about the role of properties in similarity judgements than the sentence in the text suggests. More on this in the next chapter.

¹⁵³ Fogelin does insist that the literal meaning of a metaphor and the literal comparison that is its non-elliptical counterpart are equivalent. Here the problem is the obvious non-equivalence of the literal comparison and its elliptical counterpart, so his defence of the former equivalence does not help. (It is also implausible, but that is another issue.)

An additional and serious awkwardness is the symmetry of literal 'like'-claims. In so far as (12) is true, then so is: (13) Football is like rugby. But symmetry is the last thing we would expect to hold of a metaphor; difficult as it has been to understand Romeo's claim about Juliet, the task would have been hopeless if it *required* equivalence to: The sun is Juliet. Of course, adherents of the simile view can add whatever epicycles are necessary to cope with the above problems. Thus, it could be said that, aside from its figurativeness, what makes a simile different from a literal comparison is that similes both have non-'like' equivalents—metaphors—and are not symmetrical. But these suggestions put a strain on the supposed linkage between the figurative and the literal, and they certainly weaken the defence against adhocity.¹⁵⁴

The idea of a linkage between the literal and figurative comparisons is a good one, but I think that the simile view has the dependency the wrong way around. I shall argue that literal comparison, and the notion of 'like' that goes with it, grows out of the figurative. When we see matters this way, it will be possible to see a real linkage between the literal and figurative, while at the same time preserving the spirit, if not the letter, of the simile view's idea of metaphors as elliptical similes.

An aside: strictly, this equivalence will be true only for those metaphors that are in, or can be transposed into, subject-predicate form. However, as described by the semantic descent account, any metaphor has a 'this-is-that' underlying nature, so we can see the equivalence as one which operates below the level of surface syntactic form. As already noted, the traditional simile view is not entitled to this looser construal, since its only resource for explaining metaphor is the equivalence at the surface syntactic level of metaphors and relevant similes.

4.5.4. Not returning from the dead

The argument begins where my remarks on dead metaphor left off: with the idea that the processes of metaphor figure in the background of perfectly ordinary predicates. Surprising as it might at first seem, the result of this is that we should count ordinary predicates as belonging to the class, not of dead metaphors, but to the class of expressions like 'reflect'. Fairly obviously, the latter have origins in metaphor processes: the cognitive sense of 'reflect' no doubt arose from its optical sense. But we neither need to, nor do, advert to these processes in understanding this expression, and it is most certainly not a dead metaphor. No use, however ingenious, would force us to dredge up its metaphorical origins.

¹⁵⁴ Fogelin doesn't say explicitly that the problem of symmetry threatens the linkage he assumes between literal and figurative comparison, but he must have this in mind in his heroic attempt, in the face of so much evidence to the contrary, to argue that literal comparisons are not symmetrical.

Ordinary predicate expressions like ‘rugby’, ‘football’, and ‘ewe’ are not *exactly* like ‘reflect’. The only metaphor process that figures with regard to the former expressions is qualification. This is because my story is about the *origin* of ordinary predicates. Thus, one can imagine semantic descent from the *word* ‘reflect’ to some concrete optical event, and then the use of that event to qualify a certain kind of thought process. The upshot then would be the use of ‘reflect’ to characterize a certain kind of thinking, though we wouldn't actually need to advert to this in understanding the use of this expression in the mental sphere. But, in my story about ordinary predicates, they develop directly from the use of certain objects as qualifiers of others. This beginning places the relevant qualifications outside of language. Since there is nothing linguistic for semantic descent to operate on, there is no real sense in which we are dealing even with candidates for metaphorical status.

That said, once a predicate expression has developed, there is nothing to prevent our thinking about the possibility of descent from it to an object, and, consequent on this, the possibility of using that object as a qualifier. There is thus a way, one already canvassed in a different context, in which we can after all make sense of: (12') Rugby is football, by understanding it as: (12'') Rugby is ↓football↓, in which a game of football by descent from ‘football’ can be thought of as qualifying rugby. Moreover, this possibility dovetails nicely with my treatment of ‘like’, since I am committed to regarding (12') as equivalent to: (12) Rugby is like football. Interpreting the above three sentences in the way shown allows us to preserve the linkage between ‘like’ as used in similes and its use in straightforward ‘comparisons’ like (12). In both circumstances, ‘like’ is treated as an explicit marker of semantic descent and qualification. I will return to this point, and to the issue of symmetry, shortly but the story is not quite finished.

While it is possible to see the above sentences as calling on semantic descent and qualification—these metaphor processes make sense here—we neither need to, nor do, call on these processes when dealing with these sentences. In this respect, ‘rugby’ and ‘football’ are similar to ‘reflect’: they have senses which speakers of English have mastered, and nothing requires us to take the route of descent and qualification in making them intelligible. For this reason they are not to be counted as dead metaphors, though their origin owes something to one or both of the processes of metaphor.

In the book so far, I have been studiously neutral about the best philosophical account of the senses or meanings of ordinary words, and the present issue offers no reason to change that policy. But it will help if we assume, for example, that the way to think about the sense of a predicate expression is as something like a principle (possibly principles) or property (possibly properties) by which the extension of the relevant expression is determined. Any other proposal for sense or meaning here

could serve as well, but the exposition will be clearer if we have a determinate account in mind. On this somewhat Fregean assumption, the sense of, for example, 'football' is a principle or property which determines which activities fall under this word. Thus, when someone says: 'The activity for this afternoon is football.' This sentence will count as true if the definite description picks out an activity which falls under the extension determined by the principle or property associated with 'football'.

Given a story along these lines, questions are raised when we encounter: (12) Rugby is like football. My suggestion for 'like' is that it is an explicit marker of semantic descent, but though such descent is intelligible, the words in this wholly literal sentence do not call for it. We understand them by having learnt the principles or properties associated with each of 'rugby' and 'football', and descent seems simply unnecessary. Blocked from its original role, 'like' functions in a way that is a natural extension of that role. We could put it this way: in (12) 'like' explicitly marks the way in which the sense-property of 'football' gives us information about what falls under the extension-determining property of 'rugby'. One could say that there is a kind of descent here, but to properties, and a kind of qualification of one property by another. However, while this way of putting it shows the new use of 'like' to be an extension of the original, it is admittedly a little awkward. What is it for one extension-determining property to give information about another? A simple answer in the same spirit as the original proposal is: it is for items in the extension of the subject term to fall under (at least to some extent, or in some way) the extension-determining property of the second. In terms of the example: it is for games of rugby to fall under (to some extent, or in some way) the extension-determining property or properties associated with 'football'.

This proposal for (12) has many useful consequences. To begin with, it makes intelligible, and even vindicates, the near universal temptation to see this sentence as claiming resemblance or similarity between rugby and football. This is because we think of resemblance as a sharing of a property or properties and this is certainly something that follows from my proposal. Thus, to the extent that games of rugby fall under the sense-determining property or properties of 'football', they certainly share some property or properties with football. But viewed in the way I suggest, we can make sense of various features of 'like'-claims that the mere resemblance view cannot manage easily, or cannot manage.

For instance, recall the hedging function that I attributed to 'like' in figurative claims. Counted as an explicit marker of semantic descent, it serves to alert us to juxtapositions that might otherwise be too stark to be intelligible. And it serves more or less the same function in sentences like (12). Obviously enough, rugby is not football—games of rugby are not in the proper extension of 'football'. However, when we are told that rugby is 'like' football, we are warned not to expect that an

instance of the one counts in the extension of the other; we are encouraged to think only that rugby falls under the extension-determining properties of football *to some extent or in some way*. Together with my general proposal for 'like', this also explains something simply passed over by the resemblance interpretation that goes with the traditional simile account, namely the fact that, while there is an equivalence between similes and their elliptical counterparts (metaphors), literal 'like'-claims have no counterpart elliptical equivalents. As just noted, it seems to make little sense to say that rugby is football, and therefore this cannot be equivalent to saying, perfectly sensibly, that rugby is like football. However, my proposal can explain both why the elliptical sentence seems to make so little sense and yet why, in spite of this, there is reason to regard it as *au fond* equivalent to its non-elliptical counterpart.

Recall that, by my lights:(12") Rugby is football,could be allowed to call on semantic descent. Thus, we could think of it as:(12') Rugby is ↓football↓,and, hedging aside, this comes out as equivalent to:(12) Rugby is like football.But of course semantic descent, while intelligible, is not called for in this case because we have a perfectly serviceable meaning for 'football' that simply doesn't require it. For this reason, when we first come across (12") we are bound to find it difficult to interpret as other than just plain false, and would balk at its equivalence to (12). What I am claiming then is this: we can on the one hand—and taking into account my story about the origin of predicate expressions—understand the equivalence between literal 'like'-claims and their elliptical counterparts. But, on the other hand, we would expect, and do indeed find, that the latter seem bizarrely false. In a way, we can in this case have our cake and still eat it.

I do realize that my proposal for 'like' depends on my story about the origin of predicate expressions, as well as on the apparatus of semantic descent. However, there is some independent evidence that we use 'like' in the way I have described, and not merely as synonymous with 'is similar to'. As you will see, the evidence calls on highly colloquial ways of speaking, but I don't think that makes it unreliable. In any case, even if you don't regard it as direct evidence, a brief description of these forms of speech might help to make my own proposal clearer.

Assuming that 'Ralph' refers to an adult male human being, the sentence:(14) Ralph is a man,is a straightforwardly true predication. But consider the sentence:(14s) Ralph is like a man.I take it that this sentence is not straightforward, and it requires some thought to come up with circumstances appropriate to its utterance. For example, Ralph might

be a young child, and the sentence then would be understood as making a claim sort of converse to that made about Tolstoy in (T): the idea would perhaps be that Ralph is mature beyond his years. Understood in this way, (14s) is a simile, and the claim it makes is figurative.

I will return to (14s) below, but consider now the sentence: (14c) Ralph is like: *a man*. Punctuation is necessary here (and could be achieved in lots of different ways) because the sentence as written is intended to capture a highly colloquial form of speech. Imagine that (14c) is said to a close friend by a woman who is exasperated with her boyfriend's obsession with Grand Prix racing—one among many such obsessions—to the exclusion of any activities they might properly share. The colon (it could be a comma or an em-rule) indicates a slight pause after 'like' and, as the italics show, it is then followed by an emphatic 'a man'.

Assuming, as I have been, that Ralph is indeed a man, what is going on in this and in the myriad similar uses of 'like' in common speech?¹⁵⁵ The answer, given my previous discussion, is that speakers use 'like', with its attendant pauses and emphases, as a way of forcing what by my lights is semantic descent and qualification. As I have suggested, even (14) allows for this, but we do not naturally take up this option in such an ordinary case of predication. After all, in saying that Ralph is a man, we need look no further than the simple attribution to Ralph of the properties that constitute the sense or meaning of 'man'. However, by adding 'like' and its apparatus as in (14c) to (14), speakers can overcome this reading. In forcing us to understand there to be descent to an exemplar of 'man' that is used as a qualifier, attention is focused on the properties of 'man' in a way not apparent in the simpler (14). However, since these properties do apply to Ralph directly, we do not really have a metaphor here, just the form of one that, in this case, serves a rhetorical purpose.

What then do we say about (14s)? Obviously enough, this could be a poorly transcribed version of (14c). But there are other possibilities. Taking it without the apparatus that would mark it as (14c), and being contextually discouraged from seeing it as a straightforward comparison, we could interpret it in terms of semantic descent. This, as we have seen, is most natural when we know, for example, that Ralph is a young boy, and the utterance is metaphorical. However, there is still another possibility. If Ralph was known to be an adult human male, then it is likely that the hedging function of 'like' would press itself on us. The sentence (14s) would then be taken to indicate that Ralph, though a man, is only like a man. One can imagine the negative light which a speaker of (14s) might thereby want to cast on Ralph, but, without entering into these suspect waters, it does offer further evidence that 'like' has functions other than that of straightforward comparison.

¹⁵⁵ I have heard older speakers of received English say that the widespread use of 'like' is little more than a temporizing tic of speech—something like 'uh'—which has become fashionable amongst the young. I don't think this could be right. Even without a systematic study, it is obvious enough that 'like' is used in a fixed kind of grammatical slot, and not just anywhere in a sentence. Anyway, reserve judgement until I finish the story.

4.5.5. Summary and symmetry

I return now to the main line of argument.

Any plausible version of the simile view will have to make room for figurative, in addition to literal, comparisons. This is partly because similes do not come out well if treated as straightforward comparisons of the screw/nail, rugby/football variety. But mostly it is because this is the only hope for maintaining that metaphors are elliptical similes, and of course this is at the heart of the simile view of metaphors.

Without discussion, it tends to be assumed that these two kinds of comparison are linked: that the function of 'like' is given by its role in literal comparison judgements, where it is said to mean 'is similar to' or 'shares salient properties with', and that figurative comparisons are extensions or elaborations of this role. In this way, the figurative depends on the literal. However, as we have seen, there are serious problems with this: literal 'like' judgements do not have elliptical counterparts, and the sense in which figurative comparisons involve similarity or property-sharing, as these are ordinarily understood, is stretched beyond reasonable limits. Additionally, only someone in the grip of a theory could overlook the fact that literal comparison judgements are symmetrical, whereas figurative comparisons and their elliptical counterparts (metaphors) are most certainly not.

The semantic descent account shares a commitment to a linkage between literal and figurative 'like'-claims, but, if anything, it sees the dependency the other way around. In addition to having an important hedging function, 'like' is understood as an explicit marker of semantic descent and qualification, and these are of course the very processes that figure in metaphor. So, the core understanding of 'like' involves the figurative.

An immediate consequence of this is the equivalence of simile and metaphor, and with it a sort of vindication of the simile view and an explanation of its perennial appeal. A metaphor—at least one in subject-predicate form—is an implicit simile because it makes the same call on semantic descent and qualification as its counterpart simile. And, the other way around, a simile counts as saying the same thing as its counterpart metaphor precisely because the simile explicitly calls on the same processes as its counterpart metaphor. A further consequence is that we have a natural way of understanding the hedging function of 'like': the very explicitness alerts the hearer to the processes of metaphor, and thus short-circuits problems of interpretation that might otherwise arise.

Given my claim that the core understanding of 'like' begins in the figurative, the burden for me is saying something adequate about the literal employment of that expression. I won't repeat here the details of how I think that burden should be discharged, but in outline it goes as follows. Literal 'like'-claims employ predicate expressions for which semantic descent is an option, but not one needed for intelligibility. Though these expressions have the metaphor process of qualification in their backgrounds, what counts in our current transactions with these expressions is, for example, the principle or property which determine their extensions.

The effect of this is to make the literal employment of ‘like’ only a version of the figurative; it is, in a sense, a complication. Instead of there being one object qualifying another—and thereby providing information about it—we have a property providing information about some object, and, this being something properties ordinarily do, it is easy to see why we might too hastily think of the literal use of ‘like’ as basic. When we say that: (12) Rugby is like football, the property of *being football* gives us information about the game of rugby. Of course, a game of rugby does not strictly fall into the extension of ‘football’, so (12) will not be understood as an explicit version of the straightforward predication: Rugby is football. But (12) still mirrors the figurative use of ‘like’: there is an explicit instruction to treat the predicate expression as providing information about the subject. The difference is that, in the figurative, it is the object got by descent which provides the information, whereas in the case of (12) it is the property associated with ‘football’.

A final consequence of this way of seeing the linkage between the two contexts in which ‘like’ is used is that it explains the symmetry we find in literal cases and the sharp lack of symmetry in the metaphorical. In the latter, lack of symmetry is just what one would expect, given that ‘like’ instructs us to use semantic descent on the predicate to single out an object and then use that object informatively in relation to the subject. There is no reason to think that because one object can be drafted in as a qualifier of another that the reverse should do the same, or even that it is intelligible. Thus, going along with Romeo, we can count the sun as fulfilling a predicative role in respect of Juliet, but we would scarcely expect Juliet to provide information about the sun. In contrast, in the literal case it is a *property* of one thing which provides information about another, so it is unsurprising that the reverse is likely to be true. In so far as the property of being a game of football informs us about rugby, we should rather expect that the property of being rugby should be no less apposite in informing us about the game of football. To be sure, we do not strictly have asymmetry in the one case, and symmetry in the other. The falsity of the converse doesn't follow *logically* from a figurative ‘like’-claim, nor does its truth follow in the literal case. But the semantic descent account explains why we have the expectations that we do about symmetry.¹⁵⁶

¹⁵⁶ If ‘like’ meant simply ‘is similar to’ then symmetry would indeed be a logical truth, and that is why Fogelin's insistence on the lack of symmetry in literal comparisons seems so desperate. But on my account, we can allow that there could indeed be cases in which the converse of a literal ‘like’-claim was false, though this would certainly not be the norm.

4.6. Metaphor and its Family Relations

4.6.1. 'Is a metaphor for'

The coinage is surprisingly old, but there can be no doubting the currency of the phrase 'is a metaphor for'.¹⁵⁷ The following are examples of the phenomenon:

- (i) The sunken tanker [which sank off Spain in 2002] and its unpredictable cargo which might devastate the coast at any time is a metaphor for the terrorist menace facing Western nations. (From a BBC radio broadcast)
- (ii) In retrospect, [the film] *Casablanca* can also be seen as a metaphor for Hollywood's own heroic contribution to the war effort. (From a *Guardian* newspaper article on Hollywood and its reaction to 9/11)
- (iii) All these designations capture in some imperfect way what the universe is about. It is not a clockwork mechanism, or an information processor, but it does have mechanistic and informational properties. Living organisms have goals and purposes, and I see no reason why we may not use the organism as a metaphor for the universe, as did Aristotle two and a half millenia ago. (Paul Davies writing in the *Guardian*)
- (iv) The movie [*The Hours*] is about being pinned down by social conventions and familial obligations, creating structures to make them bearable and thinking of breaking free from these fetters and liberating others ... To emphasise this, the movie begins and ends with Woolf's suicide. Before going into the water, she puts large stones in her pockets to keep her under. This is a powerful metaphor. By adding to the forces that pull her down, she ensures her liberation through death from her mental afflictions and frees her husband Leonard from the pain of caring for her. (From an *Observer* film review by Philip French)
- (v) These are fine juxtapositions, prompting all sorts of thoughts about commemoration, idolatry and the power of images over people. But the sightlines can also be wonderfully random. With the cunning use of mirrors and plate glass, you can be looking at an African coffin shaped like a Chevrolet while also catching sight of a Japanese watercolour or one of those ferocious nail figures from the Congo. Which is, presumably, a metaphor for the unpredictable movements of memory, splintering in so many different directions all at once. (From a review by Laura Cummings of an exhibition, 'The Museum of the Mind', at the British Museum in 2003)
- (vi) Her dancing [said of Isadora Duncan] was a metaphor for the way she led her life. (From a newspaper article)

These examples of the use of 'is a metaphor for' offer satisfyingly direct evidence for the semantic descent account, as well as for my treatment of 'like'. In each example, an object—where this includes, as always, events, states of affairs, etc.—is singled out by some description, and this object is then said to be 'a metaphor' for some other object. It is difficult to think of what could be going on here except in terms

¹⁵⁷ The *Oxford English Dictionary* cites 1881 as the earliest use, but I found it in Emerson's *Nature* (1846).

of the notion of qualification, though of course not under that label. Still, while some of the examples straightforwardly fit my account, there are interesting complications in others.

Beginning with the straightforward (i): the oil tanker that sank off the Atlantic coast of Spain is a wholly determinate, unfortunately actual, object. This object is described in one way in (i) rather than in any of the many other ways that would have been possible, and I shall return to this shortly. But note that it is the object itself, not its description, that is said to be ‘a metaphor for’. With only a little damage to English, one could render (i) as: The sunken tanker and its unpredictable cargo which might devastate the coast at any time *metaphors* the terrorist menace facing Western nations, where the transposition from ‘is a metaphor for’ to the more direct verb form highlights the fact that the sunken tanker itself plays the central role. What is the role that ‘metaphors’ demands of the sunken tanker? It is precisely the role that I have carved out for ‘qualifies’: the tanker gives us information of a predicative kind about the terrorist menace. Indeed, at one time I thought that ‘metaphors’ would have served instead of ‘qualifies’, but I thought better of this. The fact is that qualification is at most an *ingredient* in metaphor, and we only have metaphor proper when, in addition, there is the characteristic move from words to objects that I have called semantic descent. For all that it is an object which is said to be a metaphor for another, in each example, these objects are brought to our attention by words. The sunken tanker comes into the discussion via a description, one which not only leads us to an object, but prepares us for thinking about it in a way appropriate to the metaphor. The description of the tanker in (i) is designed to attune us to the predicative use that is made of the tanker and its cargo. The original sentence could have been put this way: The sunken tanker is a metaphor for the terrorist menace facing Western nations, and this would have brought the same object into the discussion, and given it the same role. But this abbreviated form of words would have made attunement less certain than the fuller description actually given. (The description actually given, it must be said, is so directing as to rob the metaphor of any real force. While this doesn't affect the point at issue here, I will return to it in the final section of this chapter, where I face up directly to non-syntactic richness in metaphor.)

Finally, note that ‘is a metaphor for’ consorts with ‘like’ in a way predicted by my account. As argued in the previous section, I think of ‘like’ as an explicit marker of both the move from words to an object and the use of that object as a qualifier. (As always, it is not my contention that speakers and hearers actually think of these processes under these labels and descriptions.) Thus my account has it that in this case:

Richard is like a long-playing record stuck in a groove, 'Richard' just picks out Richard, while 'the long-playing record stuck in a groove' picks out an object that plays the role of a non-linguistic predicate. Now, in the examples which opened this section, the metaphor-subject occurs *after* the object-predicate, and this means that we cannot regard 'is a metaphor' for as directly replaceable by 'like'. However, given the possibility of treating 'is a metaphor for' as 'metaphors'—a verb which aims forward to the metaphor-subject—there is nothing to stop us thinking of the passive form of this verb as equivalent to the 'like' of simile. Thus, we first have: The terrorist menace facing Western nations is *metaphored by* the sunken tanker and its unpredictable cargo which might devastate the coast at any time. And then we get a perfectly apt simile by substituting 'like' for 'metaphored by': The terrorist menace facing Western nations is *like* the sunken tanker and its unpredictable cargo which might devastate the coast at any time. Moreover, while I can understand the traditional simile theorist being tempted to claim that what is being said is that the sunken tanker and the terrorist menace are similar, we should resist this temptation. A sunken tanker and a terrorist menace are radically different kinds of thing, and the appeal to similarity here is particularly lame. But if we think of 'like' as requiring us to treat the tanker as a proto-predicate of the menace, we are not forced to explain what is going on by similarity. Since we do not regard any ordinary subject-predicate sentence as claiming that the subject is similar to the predicate, why should we think it would work in the tanker case? The tanker is not a linguistic predicate, but it serves a predicative function, and this function cannot be equated with what would anyway be a wildly implausible claim of similarity.¹⁵⁸ (Another nice feature of the equivalence of 'like' and 'metaphored by' is that the latter is, as we would expect, non-symmetric.)

So, taking example (i) as the model, my account offers a simple and direct treatment of the widely used phrase 'is a metaphor', a treatment which dovetails nicely with the story I have told about 'like' and simile. Moreover, it is a useful pointer to two further important issues.

The first revolves around the question, first raised in my introduction, of whether 'metaphor' labels a well-behaved theoretical kind, a phenomenon with an orderly underlying nature? My suggestion has been that it is such a phenomenon. The second concerns the possibility of there being a kind of semantic descent that is non-linguistic, a possibility which would open up the scope of metaphor without undermining its claim to be a kind. Each of these merits a subsection of their own,

¹⁵⁸ Of course, once we understand the claim in the way I suggest, we can see why there is a temptation to think of the tanker and the menace as similar in some special figurative way; this is at the heart of my explanation of the perennial appeal of the traditional simile account. But the mere fact that, by my lights, we have two objects present in a subject-predicate metaphor, should not fool us into thinking that similarity does any real work. The simple fact is that one of the objects serves as a predicate, and it is through that role that we understand the metaphor and not through any supposed kind of figurative similarity that the predicate object might bear to the one that is the subject of the metaphor.

in the course of which I will both explain their connection to ‘is a metaphor for’ and make use of examples (ii)–(vi).

4.6.2. Metaphor as a kind

In traditional accounts of rhetoric or style, there are lots of different families of items which are seen as making a contribution to the vividness, grace, brilliance, and energy of some piece of writing or speech.¹⁵⁹ One of these families is generally known as ‘tropes’, from the Greek for turning or conversion, and the traditional understanding here is that a trope is a device in which some word or phrase is used in a sense other than its usual one—it has been, so to speak, ‘turned’—so as to achieve one or more of the above-mentioned effects. Understood in this way, the trope family has an extremely elastic membership. Thus, while it certainly includes metaphor, it can also include such things as simile, metonymy, synecdoche, litotes, meiosis, catachresis, proverbs, hyperbole, allegory, irony, and any number of further items, depending on your taste for making ever finer distinctions. Given this variety, the idea that a trope is a word or phrase used in a sense other than the one it ‘usually’ has is not adequate as a way of delimiting a well-defined, theoretically interesting phenomenon. And certainly nothing changes here when one follows the more contemporary practice of using ‘metaphorical’ in a more liberal sense to stand in for the tropes, or for the figurative, generally.¹⁶⁰ If the metaphorical is simply counted as the use of a word or phrase in a sense other than the one it normally has, then the scepticism about my philosophical interest in metaphor mentioned in the Introduction would be justified: metaphor would be a rag-bag, studiable on a case-by-case basis, interesting in literary investigations, but not worthy of serious semantic theorizing.

Now I am never quite sure exactly what serious semantic theorizing is, whether, that is, ‘serious’ must mean ‘scientific’ in some stringent sense of this term, or something else. But it does seem to me that I have found a certain overlooked but genuine phenomenon—the use of objects, not merely as referential devices but also as predicates—and that this phenomenon can provide the basis for an account of metaphor that shows it to be a fully fledged theoretical kind.

‘Metaphor’ characterizes those linguistic acts which are rendered intelligible first, by semantic descent—a move from words to objects—and then by the employment of the resulting objects in a role that I have called, at various places, qualification, object-predication, or proto-predication. Descent and qualification are at least as robust as the processes of linguistic reference and predication they resemble, and they therefore justify metaphor's claim to be a phenomenon revealed in carving nature at its joints. Appeal to these explanatory processes has already helped us distinguish dead metaphors—expressions which, as the pun has it, are *still* metaphors—from

¹⁵⁹ This description is taken from something Suhamy 1981 notes in his survey of the figures.

¹⁶⁰ Cooper (1986: 15) suggests that the liberal use of this term actually has its origins in Aristotle, and in his interesting first chapter he offers various reasons for its revival. But nothing he says about the ‘demarcation’ problem raised by this use suggests that he believes metaphor to be in fact a theoretically unified kind.

various others that are properly ex-metaphor, expressions such as idioms and other words or phrases whose claim to metaphoricality is merely etymological. Further, many of the items on the traditional list of tropes, though they are still tropes, are easily distinguished from metaphor, when it is understood as requiring descent-plus-qualification. For example, irony clearly involves the ‘turning’ of words, but, even though the turning is not perhaps as simple as some would have it, irony is certainly intelligible without descent and qualification. So is litotes and its twin, hyperbole, as well as catachresis.

Basically, metaphor is a *special kind* of turning of words, and the semantic descent account explains both why it is special, and how this distinguishes it from many of the other members of the trope family. That this should be possible rather confirms my hunch that the phrase ‘is a metaphor for’ is actually rather important. Trendy though it may be, it is a rare first-order use of ‘metaphor’ in the vocabulary of the ordinary folk; a use which calls on the phenomenon itself, rather than merely mentioning it using a label derived from one of the more or less traditional schemes of classification. My guess is that the phrase’s being so congenial to the semantic descent account is direct, even if not certain, evidence both for the account and for the existence of metaphor as a kind.

Does all this mean that I regard a liberal use of ‘metaphor’—a use which treats it as co-extensive with ‘figurative’—as simply wrong? Perhaps unhelpfully, I think that the answer is ‘yes and no’; it depends on what you mean by figurative. If by ‘figurative’ you mean ‘non-literal’, then this would push my liberal instincts way beyond their limits. I have often claimed that metaphor involves the non-literal, but the non-literal extends far beyond the metaphorical. To take two instances just noted, irony and hyperbole are non-literal, but in my view are certainly not metaphorical. Still this leaves room for a certain liberality: there are items in the trope family, other than the traditionally and narrowly characterized trope of metaphor, that can be handled by the semantic descent account. (I am about to describe them.) So, there can be a use of ‘metaphor’ which, in including these items, is liberal. However, one must be careful about this liberality when it comes to the figurative, even when this latter is clearly distinguished from non-literality generally. This is because I believe that the figurative forms a wider class than the metaphorical, even when the latter is taken liberally. As will be discussed in the section 4.7, there are items which are certainly figurative, but are only mistakenly thought of as metaphors. Still, in view of the fact that ‘metaphor’ covers most of what is also called ‘figurative’, there is little harm in using the terms more or less interchangeably.

4.6.3. Synecdoche and metonymy

Traditionalists regarded the tropes of synecdoche and metonymy as different from metaphor, whereas the more contemporary liberal tends to include them, along with lots of other members of the trope family, in the realm of metaphor. I regard them as an interesting test case. Unlike many other members of the trope family, I do think that they belong with metaphor. There was a hint of this in my discussion of dead metaphor in section 4.4.5. Also, I think that the semantic descent account can justify this, while at the same time showing how

full-out metaphor and these tropes differ, and that, in this way, it justifies both classificatory instincts.¹⁶¹

There have always been problems in telling synecdoche and metonymy apart, and perhaps it is not worth trying too hard to do so. Traditionally, synecdoche was taken to call on the part/whole relationship as in: (15) All hands on deck (*hands are parts of crew members*), whereas metonymy was thought to call on other broader relationships, for example: (16) The White House conceded that there were no weapons of mass destruction in Iraq (*The White House is where the President lives*). But if we are willing to let 'relationship' include part/whole, as well as other less restricted linkages, then we can unify these two tropes without much loss. I shall do so, referring to the union as 'S&M'.

Unlike irony and hyperbole, S&M can stake a genuine claim to metaphoricality, and this is fully supported by the semantic descent account. However, while supporting the metaphoricality claims of S&M, the semantic descent account can also show why any instance of this trope is not exactly a metaphor. I know this sounds like having one's cake and eating it, but all will become clear shortly.

The place to begin is with an example that has nothing at all metaphorical about it: (17) The tree on the left is an oak. Leaving any controversies about descriptions on one side, there seems no harm in the naïve thought that, in (17), a description ('The tree on the left') is used to pick out an object, and a predicate ('is an oak') then characterizes that object. Bearing this simple thought in mind, consider again: (16) The White House conceded that there were no weapons of mass destruction in Iraq. Following the model of (17), it ought to be true that 'The White House' picks out an object, which is then characterized by the predicate 'conceded that there were no weapons of mass destruction in Iraq'. (Nothing in our discussion will hinge on the fact that (16) uses a descriptive name rather than a definite description.) But obviously when we treat the descriptive name this way we end up with a peculiar sentence: the object picked out by the descriptive name simply doesn't go with the predicate, since houses, even big white ones, don't make concessions.

Of course, the 'peculiarity' of this sentence is of such a familiar sort that we think nothing of it: 'The White House' refers to the President of the US, and the latter has only political difficulty in making concessions. That said, there is a

¹⁶¹ Cooper 1986 makes an interesting observation about the liberal use of 'metaphor'. He notes that using the term liberally one might say 'metaphor etc.' but never 'metonymy etc.' This is puzzling in so far as metonymy and metaphor are both in the same family, and therefore each is a reasonable candidate to represent that family. But, when you have heard my story about metonymy and semantic descent, I think the puzzle will disappear.

point in spelling out what underlies this familiar thought. And it is here that sentence (17) can help.

Keeping things at a fairly naïve level, reference is achieved in (17) by description: an object is picked out, in a relevant context, by virtue of its falling uniquely under the predicate ‘tree on the left’, and the sentence then claims that this object satisfies the ‘main’ predicate in (17). In (16) a unique object is also picked out, in this case by a descriptive name, though, as it happens, this object is not appropriate to its main predicate. Of course, I don't doubt but that we find it easy to shift from The White House to the President, an object that *does* fit the main predicate. However, it is not the ease of the shift that is in question, so much as the best way to describe it. My suggestion is that we should think of this shift in the same way that we do in the simplest of cases, that is, we should think of ourselves as looking for an object that falls under the relevant predicate and which is at the same time appropriate to the main predicate. To be sure, once we get to The White House, there is no further *linguistic* predicate around, but that is no problem. As I have so often insisted, there is nothing to stop us seeing The White House itself as a proto-predicate. Only in the case of (16), and S&M generally, this proto-predicate happens not to be the main predicate in a sentence. Instead, it serves to help us locate the ‘real’ subject of (16) by mimicking the role of predicates like ‘tree on the left’ in (17). The White House—that very object—qualifies a certain human being, albeit in a wholly conventional and unsurprising way: it is the building in which the human being is known to live. In the context of (16), the ordinary referent of ‘The White House’ doesn't work, but the referent we are led to by employing the information provided by The White House does.

Thus, in this most typical case of S&M, we have semantic descent (from ‘The White House’ to The White House), and we have a predicational use of the latter (information provided by The White House leads us to a referent for the subject of the original sentence). In short, everything needed for metaphor is present: S&M is indeed *metaphorical*, though the sentence in which this instance of S&M figures is not itself a metaphor. It might be tempting to make the point this way: with S&M we don't have metaphor, but only metaphorical reference. However, this could mislead. S&M *does* employ the processes of metaphor—there is metaphor in any instance of this trope—it is simply that the use to which the metaphor is put is internal to a sentence not itself a metaphor. It would be more accurate then to speak of metaphor being put to referential use.

In straightforward cases of S&M, metaphor processes are used to effect reference, but there is a more intricate interplay between metaphor and reference in some of those examples of ‘is a metaphor for’ which I passed over in earlier discussion. The sunken tanker example (i) seems a pure case of metaphor—a case in which an object metaphors another—but consider the Isadora Duncan example:(vi) Her dancing was a metaphor for the way she led her life. Readers of the article were assumed to know about Duncan's dancing and style of life, and were in any case reminded of these early on. Hence, when we get to (vi), it

is not so much a case of the dancing informing us about her life, so much as standing for it. Note though that the phrase 'is a metaphor' manages to fit both (i) and (vi). This flexibility means that we can understand it, as in (i), to be equivalent to 'metaphors', or, as in (vi), to be equivalent to 'can metonymically stand in for', where the metaphorical processes plays a referential role.

Nor is it always easy to tell which rendering is most appropriate. The example (iii) seems one in which 'metaphors' is most appropriate. The appeal to an organism is held genuinely to inform us about the universe; we would scarcely think that an organism stands for the universe. But example (ii), the one about the film *Casablanca* and its relationship to Hollywood's war effort, is less clear. The author goes on to stress the way in which one can understand the film's central character, Rick, as an *emblem* of Hollywood's contribution to the war effort, namely, by being patriotic and making a sacrifice in spite of being typically single-minded in pursuit of profit and self-interest. But, given the familiarity of this film, and of Hollywood's reputation, (ii) seems not so much a case of our being told something by the mention of the film, as the film's being made to stand for—to be emblematic of—Hollywood's (surprising) self-sacrifices.

The other examples follow a similar pattern in fitting one or the other category: the comment about the stones in Woolf's pocket seems informative rather than emblematic, but the remark about the layout of the British Museum exhibition seems to demand our thinking of it as standing for, rather than characterizing, memory's haphazardness. (I shall put these two examples to a different but related use in the next subsection, so there is no need to expand on them here.)

The fact that there is an interplay between a referential and a predicational function in these examples, and the added fact that it can be difficult to separate them in any given case, should come as no surprise. With them, we have circled back to my discussion of Goodman in Chapter 2, a discussion which started off from the question of whether objects could take over the roles of words.

Goodman's notion of exemplification, seen by him as a type of referential relation, was found there to offer at most a hint of the ways in which objects might take on the roles normally attributed to words. Indeed, my suggestion in Chapter 2, backed up by arguments in Chapter 3, is that we can only understand exemplification properly when it is seen as combining a predicational role with its obviously referential one. In Chapter 2, I went on to describe cases in which the referential role of objects gives way entirely to the predicational. These were examples of one object qualifying another and, except for the initial cases of tailor's swatches and paint tin lids, these examples left exemplification behind. Still, the present discussion of both S&M and the natural kindness of metaphor can give us a useful perspective on the sweep from exemplification to qualification.

What I suggest is that exemplification is merely one type of case in which the processes of metaphor are called on referentially. Thus, while it is perfectly reasonable to think that Isadora's Duncan's dancing exemplifies her life, it is less clear that exemplification is the right way to describe the *Casablanca* example. The film's character Rick may indeed be *emblematic* of Hollywood's war effort, but exemplifying

and being emblematic are subtly different notions. Each can serve to set up a referential relationship: having taken (ii) and (vi) on board, one could, in a relevant context, use the dancing or the film metonymically. Think of the following continuations of (ii) and (vi):(ii-c) In the 50s, *Casablanca* was superseded by a return to mindless pursuit of box office success.(vi-c) Isadora's dance did not figure widely amongst the lifestyles of those suffering in the Depression. But, unlike a standard case of exemplification, *Casablanca* and the character Rick are not actually instances of the property or properties true of Hollywood's efforts during the war. So, it is inappropriate to describe what is going on as exemplification. (Rick is not a film-maker, nor is *Casablanca* a film about making films, but Isadora Duncan's dance is typical of the things she did, of her style of life.)

Obviously, 'exemplifies' is not a precise term, but perhaps the point I am making will be clearer if I use the original S&M examples. The White House in (16) is understood to refer to the President, and it makes perfect sense to say that this building is an emblem of the Presidency. But it sounds distinctly odd to say that the White House exemplifies either the man or the office. Similarly, though we have no trouble at all in understanding a pair of hands as referring to a human being, we would neither say that the hands exemplify nor that they are emblematic of a human being.

The cases in which the processes of metaphor serve the more mundane task of reference should be grouped together, and they range from S&M to Goodmanian exemplification and its near relatives. To be sure, there are differences, most particularly in respect of novelty. We take in common-or-garden cases of S&M such as 'all hands on deck' without a first, never mind a second, thought, but cases like (ii) and (vi), though they are not exactly powerful, require some reflection. It can be something like a discovery that *Casablanca* or a dance can exemplify or be emblematic. Still, in so far as all these cases involve semantic descent and qualification, along with S&M, they fit well with central cases of metaphor, cases in which descent and qualification explain the thrust of the whole of an utterance and not just its referential part.

The scope of 'is a metaphor for' supports this grouping, and suggests that the folk themselves think of metaphor as a genuine kind with just the extension that I have so far identified. As already noted, this phrase manages to cover *both* cases of metaphor proper, *and* cases of metaphor in the service of reference. This is evident in the examples which opened section 4.6. Nor would it be out of place if this expression were used in regard to cases of S&M: it makes perfect sense to say that The White House is a metaphor for the Presidency, or that hands are a metaphor for whole human beings in some active role. These are not in any way original or stunning, but being original or stunning is not actually a requirement of metaphor.

4.6.4. Visual descent and visual metaphor

This final subsection makes a concession. So far I have spoken of semantic descent as a move from words to objects, and

I concede in what follows that this is too restrictive. But the possibilities that this concession will open up are too vast to be pursued in any detail in this book, so what I say should only be regarded as a marker for further work.

The place to begin is with the examples of ‘is a metaphor for’ that I only touched on above. Consider again this part of one of the examples: Before going into the water, [Woolf] puts large stones in her pockets to keep her under. This is a powerful metaphor. By adding to the forces that pull her down, she ensures her liberation through death from her mental afflictions and frees her husband Leonard from the pain of caring for her. This is subtly different from the examples so far discussed, though not for the trivial reason that the exact words ‘is a metaphor for’ do not occur in it. These other examples began with descriptions of some object—an object which is then said to be a metaphor for something else—whereas in the film review, it is the image of an object that figures centrally. That is, the reviewer describes the film image of the Woolf character putting stones in her pockets, and seems to be claiming that this image is a metaphor. Now, as a matter of fact, the real Virginia Woolf was found to have put stones in her pockets before she entered the River Stour, so the film image does reflect historical reality. But it would be carrying things too far to suggest that Woolf did so for metaphorical reasons, and, in any case, the reviewer is not speaking about what she did, but about what the film and its creators did. Still, there is something right about its being the placing of the stones, rather than merely the image of this placement, that does the metaphorical work. So, what is actually being claimed as the metaphor: the placing of the stones by Woolf, or the image of that act achieved by having an actor playing Woolf putting stones in her pocket?

The solution to this mild puzzle requires me to widen my conception of semantic descent, and this is the concession mentioned above. Up to now, I have said more than once that semantic descent is a movement from words to objects. But what I think this example forces me to say, something welcome in any case, is that semantic descent can also correctly describe the move from images, depictions, and other kinds of non-linguistic representation to objects. In the example above, there are complications: we have the reviewer using *words* to describe the film *image*, and then we have the film image representing Woolf’s having put stones in her pocket. With this more liberal conception of semantic descent there are thus two possible kinds of descent in the example: from words to an image and from an image to an actual event. What I suggest is this: the metaphor consists in the film-maker using a film image to represent a certain event—Woolf putting stones in her pockets—and then using that event to qualify the profoundly disturbing situation which Woolf felt herself to be in. (Note: the qualifying event is the historical one, and so is the state of affairs it qualifies.) What then of the words of the reviewer?

I think the right construal is this: when the reviewer describes the stones being put in the Woolf character’s pockets as a powerful metaphor, he is calling attention to a metaphor—a visual metaphor—but the metaphor itself is the work of the film-maker. For this reason, there is an important difference between (iv) and, say, (i).

The author of the words in (i) is also the author of the metaphor, and, by my lights, this comes about because the phrase ‘is a metaphor for’ implicitly calls on my account to make the sentence intelligible. Thus, the author uses words to describe the sunken tanker—inviting thereby semantic descent to the sunken tanker itself—and he then claims (using ‘is a metaphor for’) that this object qualifies the terrorist menace facing Western nations.

In contrast, the author of (iv) is not the author of the metaphor he asks us to notice, nor would it be appropriate to understand him to be claiming that the *film image* is a metaphor for Woolf's situation. It is not the image, it is the event represented by that image that is the metaphor for Woolf's situation.

Of course, there is another way to read (iv). Suppose that the film-maker had not in fact intended us to take the event depicted in the way just described; suppose that, in order to give credence to the film, it was offered merely as a representation of an event known to have happened, and for no other purpose. In this circumstance, perhaps we should regard the reviewer as the source of the metaphor. Having watched the film, and seen its graphic depiction of the suicide as it might have happened, the reviewer finds the placing of the stones to be a metaphor for Woolf's terrible situation (as she saw it). This would make (iv) much more like (i), and it would give the film image a somewhat different role. The image would have prompted the reviewer to think about the historical event, but it would not have played a part in the metaphor itself.¹⁶²

To show the parity of visual and linguistic metaphor, consider example (v), which concludes this way: With the cunning use of mirrors and plateglass, you can be looking at an African coffin shaped like a Chevrolet while also catching sight of a Japanese watercolour or one of those ferocious nail figures from the Congo. Which is, presumably, a metaphor for the unpredictable movements of memory, splintering in so many different directions all at once. This seems to me also to be a case of visual metaphor, though there is no image, and the object which results from semantic descent metonymically stands for, rather than directly qualifies, memory. Let me explain.

What plays the role of an image or a depiction in (v) is the actual layout of the exhibition. While the layout might have come about inadvertently, it didn't. The designer of the exhibition used the layout to call attention to itself, or, more in keeping with the language of art exhibitions, it served a self-referential purpose. We were meant to be struck by its haphazard nature and, given the subject matter of the exhibition (‘Museum of the Mind’), we were then invited to make a connection to the faculty of memory itself. When the reviewer insists that the layout is a metaphor for the unpredictable movements of memory, this should be understood as involving metonymic (hence metaphorical) reference. The exhibition's layout refers us to memory, rather than adding to our knowledge of this faculty.

¹⁶² For what it's worth, the full text of the review, and even simply the full text of (iv) make the first interpretation the overwhelming favourite.

I do not claim that the possibility of visual metaphor is revealed here for the first time. What I would claim as novel, though, is the way in which my account shows this phenomenon to be seamlessly continuous with linguistic metaphor. Once we allow that semantic descent is not merely a word–object relationship, that it also describes the move from pictures and other representations to objects, the rest of the semantic descent account of metaphor guarantees that, when we speak of ‘visual metaphor’, this involves no stretching of ‘metaphor’. Non-linguistic representations take us to objects and, in the right context, these objects qualify others. The end result is visual metaphors which share all the properties so far attributed to their linguistically based cousins. And, among these properties, is the possibility of assertoric force.

In Chapter 2, I briefly touched on some remarks that Goodman made which cast doubt on the idea of a picture making an assertion. To remind you, he wrote: ‘I like to keep the term “true” for statements. Statements in a language are true or they are false. I don't like to speak of a picture as being true or false, since it doesn't literally make a statement.’ And he supports this claim by noting: ‘A picture like a predicate may denote certain events ... When the predicates in a text denote those same events ..., the picture and the text are to that extent intertranslatable; and the picture, though it makes no statement, might be derivatively called true or false according as the text is. But we must not forget that, strictly speaking, calling a picture true or false is false.’

What I have said about visual metaphor suggests that, while Goodman is not just plain wrong, there is something mistaken about his restriction of truth to linguistic statements. He is certainly right to say that a picture, or any other sort of representation, can be used to denote objects or events, just as a linguistic predicate can. But what Goodman seems to overlook is that the objects thereby denoted might tell a story of their own, and stories are apt for being judged true or false.

Though already implicit in what was said above, it is worth spelling out just what Goodman missed. The place to begin is with words. Assume (what may be true) that in the novel on which *The Hours* film was based, there are some words to the effect that the Woolf character puts stones in her pocket. Assume too that the words in the novel were part of some metaphor. With apologies to the author, suppose they were: ‘As she was about to walk into the swirling waters of the River Stour, she carefully put stones in her pockets. These stones were her afflictions, and they would liberate her.’ On my account, the metaphor consists of, first, semantic descent from ‘these stones’ to the historical event that is only fictionally described in the novel, and, second, the use of this event to tell us about (to qualify) Woolf's afflictions and her attitude towards them. Of course, someone might insist that what it tells us is false, that Woolf drowned herself for reasons other than those that, in the novel, are suggested by this metaphor. And this is just what we would expect: though in many contexts, we are simply not interested in the truth of what is said by a metaphor, when we are, the metaphor takes its chances just like any other assertion.

Now, Goodman allows that the words characterizing the stones-in-the-pocket event and the film image are 'intertranslatable': they each denote the same event, the one involving the historical Woolf. But, given this, and given my account of how that event comes to be used by the writer to make an assertion, can we not say precisely the same thing about the film-maker's use of the image? That is, why not regard the presentation of the historical event by the film image as part of what would then be a visual metaphor used to make an assertion? Goodman denies that this can happen, but then again he never considered the use of objects as qualifiers. Besides, we are often tempted by the thought that, even without any theoretical backing, films and other non-linguistic representations somehow do make assertions. So, an account that cashes in this 'somehow' should be appealing.

Finally, I offer a cautionary note about visual metaphor. I have throughout insisted that the phenomenon of qualification on its own is *not* metaphor. There can be exercises of our semantic attitude towards objects—cases in which we find predicative information in those objects—without these counting as metaphor. This was part of the lesson of my examples in Chapter 2. The disjointed paling that Nabokov noticed, the fallen tree seen by Smith during a storm, these were informative—in each case they served as proto-predicates of relevant subject matters—but they were not themselves metaphors. Only when there is direction, of others or ourselves, to some object, whether by words or, as is now conceded, by images, and at the same time the need to find a qualifying function in that object, do we have metaphor.¹⁶³

Given the possibility of visual metaphor, one has to be especially careful about confusing mere qualification with the use of qualification within metaphor. When dealing with certain kinds of non-linguistic representation, it is all too easy to overlook this distinction. For example, consider again the visual metaphor in the Museum of the Mind exhibition. The layout of the various items in the exhibition was said to be a metaphor for the nature of memory. But a layout of things in various rooms doesn't seem to be all that different from Nabokov's paling or Jones's fallen tree: neither of these latter are representations of anything, they are simply features of the world that we come across. So, why is the exhibition intelligible as a visual metaphor? The answer is contained in earlier remarks about the intentions of the designer. The layout of the exhibition was not merely a state of affairs which just happened to be the case, the designer intended this layout to call attention to itself. Nor was this simply a matter of the designer's wanting people to notice the care or skill with which it was done. If what was said in (v) was right, the designer wanted us to take the layout self-referentially and as a result the information contained in the layout becomes a metonymic stand-in for memory. (The layout calls attention to itself and, given its haphazard organization, it serves as a proto-predicate leading us to memory.)

¹⁶³ There should be nothing controversial about this. We would not regard my noticing that some object is made of wood as my asserting it to be made of wood. For assertion, at least of the simplest sort, I have to single out an object (refer to it) and use some device (a predicate) for conveying that the object is made of wood. The need for reference and predication in metaphor assertion is no different, though predication is accomplished not by words, but by objects.

This referentiality, and the use to which it is put, lacking in cases of mere qualification, justifies our regarding the exhibition's layout as containing visual metaphor.

The possibility that perfectly ordinary objects (including events and states of affairs) can be referential, and thereby can figure in non-linguistic metaphors, seems to me to have consequences for what is called 'conceptual' art. There simply isn't space here to consider those consequences at any length here, but the general idea is simple enough. In the 'installations' which figure so prominently in conceptual art, one has just the kind of referential intention necessary to visual metaphor. The objects so installed are not merely in front of us, they are *put* in front of us (installed) so that we will attend to them in a special way. Any object displayed in some public space demands attention, but the point about the objects installed by a conceptual artist is that we are invited to take them as having a point beyond any intrinsic interest they might have as objects on display. (Often the objects are chosen precisely because they have little intrinsic interest. This is presumably an economical way of getting us to think beyond the object.)

From the perspective developed here, this further point is that of proto-predication or qualification: the objects installed are offered as sources of information about some subject matter, other than that of the installation. For this reason, the installation can be counted as a visual metaphor, or as containing visual metaphors. However, this might be too strong, since a certain coyness or indirection is often part of the enterprise. We are unlikely to be told what the further subject matter is, or given hints about the kind of information that is available in the objects. This suggests that what the artist offers is not the object as a visual metaphor, but merely the object as an occasion for qualification. Thus, imagine an artist being somehow responsible for putting a wind-blown, fallen tree in the path of walkers. Smith happened along and took it as a source of information about her own situation; others would take it differently or not at all. It is as if the aim here is the installation of a wholly natural occurrence that might be taken as informative in some way or other, an aspiration that is intelligible, even if achieving it is in fact impossible.

Whether conceptual art is best understood as visual metaphor, or merely as an occasion for qualification that falls short of the explicitness of metaphor, the materials for understanding it seem to be precisely those that figure in the semantic descent account of metaphor. Further discussion, of necessity on a case-by-case basis, is simply not possible here.

4.7. The Figurative and the Metaphorical

In the previous section, I noted that some writers use 'figurative' to mean 'non-literal'. There is nothing wrong with this, but it seems a shame to waste the potentially useful 'figurative' as a synonym for the unruly 'non-literal'.¹⁶⁴ My suggestion was

¹⁶⁴ The *Oxford English Dictionary* offers 'abounding in figures of speech' as one meaning of 'figurative', and this rather supports the wider use of the term. But its other meanings, namely, 'metaphorical', 'emblematic', 'pictorial', are more congenial to my understanding.

that we should regard the figurative as closer to the metaphorical and, perhaps somewhat mysteriously, I claimed that we should understand the figurative as *almost* co-extensive with the metaphorical. In this section, I shall explain the 'almost'. Along the way, the discussion will take in the phenomenon of synesthesia, one which has been claimed (I think wrongly) to be at the heart of metaphor.

4.7.1. Lumpiness

The place to begin is with a feature of metaphor that tends not to be noticed, but which is interesting in its own right, and gains an easy explanation on the semantic descent account. While no label I can think of is completely appropriate, I tend to think of the phenomenon as the essential 'lumpiness' of metaphor.

For the purposes of explaining what I have in mind, nothing will be lost by sticking to the simpler subject-predicate form of metaphor: A is F. Almost without exception, the filling for 'F' is the name, often a count noun, of some determinate object (including, as always, such things as events and states of affairs), such as 'a fox', 'the sun', 'a summer's day', 'the shape of a landscape'. There are lots of what are sometimes called 'pure' predicates that could in principle fill the 'F' place, but the result of using these tends not to be a metaphor. To be sure, we do, or at least did, speak of emotional problems being 'heavy' and certain people being 'square' or 'sweet', but when we do these tend to be derivative from less elliptical expressions, namely: Amy's emotional problems were a heavy weight, or John is a Euclidean regular figure, a square, or his character has the appeal of sweet-tasting food.

That metaphors tend to be lumpy—tend to call on determinate and often concrete objects to make their point—is of course highly congenial to the semantic descent account. For obviously enough, in order for qualification to work, there must be an object that can be reached by semantic descent. If 'What John experienced yesterday was heavy' was not in fact elliptical, there would be no obvious candidate to be descended *to*, and which could then qualify John's experience.

It might be conceded that 'heavy' is a metaphorical way of describing someone's experiences only because we take it as elliptical for 'heavy physical object', a phrase describing an object difficult to carry around. But, it will be countered, there surely are 'pure' predicates which are used in metaphors without being elliptical in this way. Here are two obvious cases: The Largo from Dvořák's New World Symphony is sad, When he heard he didn't get the job, John was blue. These seem obvious candidates for what we might think of as 'smooth' metaphor; metaphor without the lumpiness associated with semantic descent to objects. Moreover, don't such cases undermine not only my claim about the essential lumpiness of metaphor, but also the support I find for the semantic descent account in this alleged feature? Indeed, don't these cases suggest that metaphor is not, after all, quite the homogeneous kind I have taken it to be?

In order to preserve the thought that metaphor can be understood as a circumscribed phenomenon which works via descent and qualification, I have to deny that the above examples are metaphors. And I certainly do deny this. Yet we are so used to thinking of them as metaphorical that mere denial is not enough: some justificatory story must be told.

The opening remarks of this section anticipate the basic idea behind that story: the above examples are figurative uses of language, by which is meant something more specific than non-literality, but they are not actually metaphors. However, the story itself requires me to begin by saying something about a strange phenomenon that has only gained a certain respectability (as a phenomenon) in the past few years.

4.7.2. Synaesthesia

A recent article discussing the phenomenon begins this way: When Mathew Blakeslee shapes hamburger patties with his hands, he experiences a vivid bitter taste in his mouth. Emerelda Jones (a pseudonym) sees blue when she listens to the note C-sharp played on the piano; other notes evoke different hues—so much so that the piano keys are actually color-coded, making it easier for her to remember and play musical scales. And when Jeff Coleman looks at printed black numbers, he sees them in color, each a different hue. Blakeslee, Jones and Coleman are among a handful of otherwise normal people who have synesthesia. They experience the ordinary world in extraordinary ways. (Ramachandran and Hubbard 2003: 43)

The authors are cognitive psychologists (of a distinctively neurological tinge) and in the article they consider experiments showing the robustness of synaesthetic phenomena, and offer an explanation of them in terms of ‘cross-activation’ (where this is not to be confused with mere ‘cross-wiring’) in the parts of the brain responsible for sensory processing. The evidence that synaesthesia is a robust phenomenon is necessary. Not only were the claims of synaesthetes once thought to be fraudulent, even those who took them seriously often regarded them as memory reports rather than reports of sensory experience. The bizarreness of the idea of someone's seeing a black numeral ‘5’ as actually red sends us rushing to the simpler hypothesis that such a person merely associates redness with 5, or perhaps with a certain numeral-style of 5, and is reporting this memory association rather than any direct visual experience. However, the experiments of Ramachandran and Hubbard, and others, convincingly rule out the ‘simpler’ hypothesis, at least in respect of visual synaesthesia. (Less convincing is the case for cross-modal synaesthesia: seeing blue when hearing C-sharp played on the piano, but this is not the place to thrash that out.¹⁶⁵)

¹⁶⁵ In correspondence, Hubbard has agreed that there are conceptual problems in respect of the cross-modal cases. The synaesthete who sees what appears to us a black ‘5’ as red is odd, but at least the colour so perceived occupies a space in the visual field—the same space as the numeral that we also see—so it makes sense to speak of his seeing the numeral as red rather than as black. (Or at least does so for some subclass of subjects; others describe the colour as merely hovering around the numeral.) But do we want to say, on analogy with this, that certain people can hear blue? Can we even make sense of this? Note that in the quote Ramachandran and Hubbard speak of ‘seeing blue when hearing C-sharp’. This suggests more of an association than a genuine case of employing a sensory modality that we regard as auditory for tasks that we regard as visual. The tests in the visual case show that the synaesthete can reliably sort what we see as a jumble of black numerals in ways that we could only do if the numerals were indeed of different colours. It is not clear whether there could be tests that had this form in cross-modal cases. It should be noted that in the fast developing study of synaesthesia, Ramachandran and Hubbard belong to the group of theorists who believe it to be a phenomenon originating at the lower levels of perceptual processing. Hence, it is not surprising to find them suggesting that synaesthetes can actually hear sounds. But much other work suggests that synaesthesia, while not less puzzling, is really something like an association of sensory modalities that is made at a higher level of cognitive processing. (Note that this kind of association is a genuine association of sensory outputs and should not be confused with the appeal to memory association made by early septsics.)

After detailing a neurological suggestion about how synaesthesia works, and having claimed that synaesthesia is seven times more common in ‘creative’ people than in the general population, Ramachandran and Hubbard offer this speculation: One skill that many creative people share is a facility for using metaphor (‘It is the east, and Juliet is the sun.’) It is as if their brains are set up to make links between seemingly unrelated domains—such as the sun and a beautiful young woman. In other words, just as synesthesia involves making arbitrary links between seemingly unrelated perceptual entities such as colors and numbers, metaphor involves making links between seemingly unrelated conceptual domains. Perhaps this is not just a coincidence. (Ramachandran and Hubbard: 47)

Leaving on one side the tendentious claim about creativity, it seems to me that this speculation involves a fundamental error, but one that is highly instructive. Whether synaesthesia involves a genuine cross-over in perceptual awareness, or involves a higher level ‘leakage’ from one perceptual output to another, it is not wrong to describe it as, in some fundamental sense, associative. Moreover, it is not wrong to describe it as *purely* associative. But metaphor is just not like this, nor is this simply a matter of degree, with metaphor at some higher cognitive level and synaesthesia further down. Metaphor is not in any sense a matter of association. What I mean will be clearer if I begin with a non-synaesthetic example of what I mean by ‘associative’.

When I reach out to touch a smooth curved surface, my awareness of the surface and its curvature take a certain form, one which differs from my awareness of a flat surface, or an irregular one. Without lifting the lid too high on this can of worms, it seems reasonable enough to say that these different awarenesses are linked to the surfaces that give rise to them by means of some pretty complicated goings-on in my nervous system. And when I say ‘linked’ here, I mean to be committed to nothing more profound than the idea that the awareness and the neurophysiology are in step, and that this is a matter of a fact about me. Now, there is lots more to discover about this linkage, even leaving on one side, as I have promised, the explanatory gap between neurology and awareness. But there is something that I can say now about the linkage which is true, and will not be overturned by any future finding: no one will ever find a reason or a justification for particular tactile awarenesses in features of the surfaces that I touch, and the nature of my awareness in each case will not itself count as a justification for a surface’s being one way or another. That curvedness feels to me as it does is an association supported by the kind of nervous system I have: it is, so far as reasons go, *just* an association.

Though the synaesthete's experience is very different from most people's, there is nothing different about the *type* of linkage that characterizes this experience. When I look at a numeral of 5 printed in black, I am aware of it as having a certain shape and colour. There is no doubt a linkage between this experience, my nervous system, and the surface reflectance of the numeral on the page, and even if it doesn't fully explain my experience, this linkage is a perfectly reasonable subject of investigation. Moreover, as explained above, it is a purely associative link: the reflectance and the neural goings-on, even if I knew all about them, would not provide me with *reasons* for the numeral appearing as it does. When the synaesthete looks at it, he too is aware of it as having a certain shape and colour, only his judgement of the colour differs radically from mine, even though, unlike someone colour-blind, the synaesthete has no problem with colour perception generally. What this suggests—something certainly worth researching—is that the linkage in place in the synaesthete's nervous system is interestingly different from that present in 'normal' perceivers. When the research is complete, we will have an even better idea of how the synaesthete differs from normal perceivers, and we will understand the special kind of linkage in the synaesthete's central nervous system. However, knowledge of this will no more provide the synaesthete with reasons or a justification for the numeral's appearing as it does, than knowledge of the linkage in normal perceivers provides us with such reasons.

Ramachandran and Hubbard hint at the point just made when they say that synaesthesia involves the brain making 'arbitrary' connections between perceptual appearances. Though this is not perhaps the best way of putting it—given the role of the central nervous system, the connections are most certainly not arbitrary—one can see what they are getting at. Only a cause, and nothing like a reason or justification, can be given for the synaesthete linking what are distinct ('seemingly unrelated') experiences in normal perceivers.

Alright, so let's suppose that I am right that the relevant neurologically grounded links in both normal and synaesthetic perception are merely associative, what has this to do with metaphor? Well, given everything that has been said about metaphor, the obvious and correct answer is 'nothing much'. That is to say, in metaphor, the link between the subject of the metaphor (e.g. Juliet) and metaphorical 'vehicle' (e.g. the sun) is very far from being merely associative. Romeo, one must imagine, thought there were good reasons for making this connection, reasons that he also thought his audience would appreciate. Just to make the obvious even more so, Romeo's reasons were of two sorts: those which led him to think that the sun was apt for the kind of information he wanted to convey, and those which led him to think that the information applied to Juliet. In my account, the first set of reasons justify the use of the sun as a proto-predicate with a certain content, and the second justify the application of that proto-predicate to Juliet. But the general point about the rational connection between metaphor subject and vehicle is common to any account.

The earlier citation from Ramachandran and Hubbard suggests that they were in some sense aware of this. The linkage in synaesthesia was described as an arbitrary

one between unrelated perceptual entities, but, in the case of metaphor, they drop ‘arbitrary’. Thus, they speculate that creative people’s ‘brains are set up to make links between seemingly unrelated domains’. But even though they do in this way implicitly distinguish between synaesthesia and metaphor, they do not see the importance of the distinction. To repeat: the simple fact is that perceptual linkages, whether normal or synaesthetic, are merely associative, while it is of the essence of metaphor that the link, however unexpected, is one it makes sense to try to justify, to give reasons for, and generally to use our rational faculty to understand. The bottom line: metaphor is not in any important sense like synaesthesia.

4.7.3. Normal perceivers and synaesthesia

In reaching the above conclusion, I took what might appear an unnecessary diversion through the case of normal perceivers. That is, I claimed that normal perception involves associative linkages between perceptual objects and experience, and to that extent is no different from synaesthetic perception. What I should like now to show is that this was not so much a diversion as it might have seemed. The joining of normal and synaesthetic perception was intended to soften you up to the idea that we can also understand certain things about normal perception by thinking of it in terms of synaesthesia. And when you see what I mean here, I think you will understand why certain common figurative claims, though they are often said to be, are not after all metaphors.

One theme in current research on synaesthesia is that the condition, bizarre though it can seem, is in fact continuous with normal perception. It has even been suggested the modes of perception in infants are sufficiently jumbled up to make synaesthetic linkages for infants the norm, but that these linkages drop away during maturation (see Rich and Mattingly 2002). Moreover, whatever turns out to be about the full neurological story of its origins, the possibility that synaesthesia could be a feature of normal perception, and could persist in some form into adulthood, shouldn’t be surprising. Indeed, there is nothing to prevent us thinking of certain forms of synaesthesia as the norm. Of course, given that the label ‘synaesthesia’ tends to be reserved for what is statistically an abnormality, it might sound odd to speak of synaesthesia as the norm. But a better way to think of synaesthesia is as the linking of distinguishable perceptual modes, and this allows for the possibility of some kinds of synaesthesia being common, and others as rare, and therefore strange.¹⁶⁶

Ramachandran and Hubbard seem to accept that some kinds of synaesthesia might be the norm. In the website that accompanied their *Scientific American* article, they write: A form of synesthesia exists in all our brains. For instance, we speak of certain smells of particular liquids—like nail polish—as being sweet, even though we have never tasted them. This might involve the close neural links and cross activations—that is, when one area of the

¹⁶⁶ The idea would be that one distinguished perceptual modes by the kind of thing they were perceptions of, so that sound-colour linking would come out as synaesthetic even if everyone made that link. This obviously raises questions about how narrowly to define the objects of perception, but these questions are not relevant here.

brain affects another—between the areas involved in smell and taste. This would make sense functionally—e.g., fruits are sweet and also smell sweet like acetone. But it also makes sense structurally, because the brain pathways for smell and taste are closely intermingled and they both project signals to the same parts of the frontal cortex during sensory processing.

Assuming, as most work on synaesthesia allows, that cross-linking may begin in the perceptual realm, though is not confined to it, my suggestion is that, when we think of certain colours or sounds as having this or that emotional effect, we are giving evidence of a form of synaesthesia. Thus, when someone says that he is blue or that the music is sad, the ultimate bases for these claims are synaesthetic linkages, ones widely enough shared to have made these kinds of claim intelligible.¹⁶⁷ Moreover, in so far as these claims are grounded on synaesthetic linkages, the earlier point that I made about metaphor and synaesthesia comes into play. To repeat what was said: in metaphor there is always the possibility of finding reasons or justifications for the relevant predications, but in synaesthesia, the connection is essentially brute. It makes no more sense to try to convince someone that a mood can be blue, or that a certain sound pattern is sad, than it does to try to convince someone by argument that, say, a standard pillar box in normal lighting is red. There is no doubt a story to be told about why human beings perceive it to be red, but the story cannot show that the pillar box *merits* this reaction. In contrast, it is always intelligible to wonder about, not only the truth, but the aptness of a metaphor, and in many cases we do this at length.

Much more could be said about the blue/sad examples, and, though a lengthy discussion isn't necessary for the point they illustrate, a couple of comments are in order. Colours are conventionally associated with lots of different things in different cultures, so I can imagine doubts about my claim of a synaesthetic connection between blue and, as one dictionary has it, being down in the dumps. Two things should be said about these doubts. First, the fact that colour words and feelings come to be linked conventionally is in no way an objection to the contrast I draw between metaphor and these linkages. For, even if it were true that the connection between blue and the dumps is purely conventional, this connection would be quite different from, say, the metaphorical connection between Juliet and the sun. But, second, I doubt that these connections could be *purely* conventional. While one shouldn't exaggerate the claim of a synaesthetic connection, I suggest that it is in some way at least minimally responsible for the blue/dumps link, whatever the actual origin of 'blue' used of mood. The colour spectrum is experienced by us as roughly divided into the hot and the cool, and the spectrum of human emotional

¹⁶⁷ The actual origin of 'blue' for mood in this case is controversial. However, my claim doesn't depend on its origin being exactly as I have described. For example, it needn't be true that these kind of linkages are universally shared: as so often happens, some subsection of the human community might be a source for the 'blue–depressed' link which then becomes purely conventional for the rest of us. Or, as will be discussed in the text, it might be that there is a wholly different origin for the use of 'blue' as of mood, but it sticks because of an associative link. Note that the sad-music case merits more discussion because of the social dimension of emotions like sadness, but I do not think that any elaboration would undermine the point made in the text.

states also tends to be conceived as having a bipolarity. All that is necessary for there to be a synaesthetic basis to what are now conventional links is that there be some linkage between these two divisions. While I stand ready to be surprised, it would be surprising to find that some culture used ‘red’ as the conventional marker of someone’s feeling in the dumps.¹⁶⁸ Moreover, the pairing need not be actually experienced each time the words are called on; so long as most human beings would not find it bizarre to speak of *feeling blue*, when they are down in the dumps, convention suffices to maintain the currency of this talk

Speaking of music as sad is much more complex, and even the little that I have claimed is liable to misunderstanding. It might be thought, for example, that my claim of synaesthesia comes to no more than the familiar (and mistaken) idea that ‘The music is sad’ means ‘The music makes me feel sad.’ But this is not the case. While it is true that, because of the synaesthetic connection between certain sounds and feelings, someone may actually be cast down by the sounds, this is certainly not what it means to claim that those sounds are sad. What the latter claim asserts is that certain sounds, like certain feelings, can be perfectly well described as sad. And synesthesia grounds the linkage between the sounds and the feelings that makes this sharing of a predicate intelligible.

Returning to the main point: I have argued that pure predicative claims (e.g. ‘John is blue’) are not metaphors. They, therefore, do not threaten either my inclination to regard metaphor as ‘lumpy’, or my insistence that it involves essentially both semantic descent and qualification. That said, it is perfectly reasonable to think of these pure predications, as we do anyway, as figurative. Among its meanings, ‘figurative’ suggests the idea of an emblem—think here of ‘figure-head’—and it is perfectly reasonable to think that ‘blue’ or ‘sad’ could come to be emblematic of certain kinds of feeling or music. They can be thought of as leading us to a kind of feeling or a kind of sound. In so doing, these words function something like metonyms, but with this crucial difference: the route from a genuine metonym to a referent goes via proto-predication: the metonymic object contains information about the object it leads us to. In contrast, ‘blue’ (said of a person) and ‘sad’ (said of some stretch of sound) do not offer an informational route to a kind of feeling or kind of music; there is instead a direct linkage set up by the appropriateness of our shared synaesthetic reactions, and this is kept in place by convention.

One final note that I think clears up something puzzling in Wittgenstein, and, more pertinently, reinforces the line I have pursued in this section. He wrote: ‘The secondary sense is not a “metaphorical” sense. If I say “For me the vowel *e* is yellow” I do not mean: “yellow” in a metaphorical sense,—for I could not express what I want to say in any other way than by means of the idea of “yellow”’ (Wittgenstein 1953: 216). This passage comes as a sort of coda to his reflections on: Given the two ideas ‘fat’ and ‘lean’, would you be rather inclined to say that Wednesday was fat and Tuesday was lean, or the other way round? (I incline to the former.) ... Now I say

¹⁶⁸ Holly Golightly in *Breakfast at Tiffany’s* speaks of a sort of depression as her having the ‘mean reds’. But presumably the point of this is to underline how different it is from melancholy.

nothing about the causes of this phenomenon. They might be associations from my childhood. But that is a hypothesis. Whatever the explanation,—the inclination is there. (Wittgenstein 1953: 216)

What I think Wittgenstein claims is that precisely because we make a purely associative connection between Wednesday and ‘fat’, we should *not* say that it has a metaphorical sense. That is, although we know that ‘fat’ doesn’t include Wednesday in its extension—it is not part of its ‘primary sense’—this is no reason to say that it has, as he puts it, a ‘secondary’, metaphorical sense. The example of the letter *e* most naturally brings to mind synaesthesia, though this is not explicit. But what he says about ‘fat’ would apply equally to synaesthesia: like a childhood association, a synaesthetic link is just an association. It is not something for which we can give a reason. When therefore he claims (in effect): just because we cannot express the point without ‘yellow’, we should not therefore think of ‘yellow’ as metaphorical, he is noting the fact that metaphor requires something more than a purely associative link. And, though he does not here say what would be required, I find his sentiment (and the tenor of the thoughts around the above passages) wholly congenial to my line of argument about synaesthesia and the figurative.

4.8. The Riches of Metaphor

In section 4.2 of this chapter, I described the way in which my account accommodates syntactically complex metaphors. In the remaining sections, I have considered: mixed, dead, extended metaphors; the relationship between simile and metaphor; a certain common expression ‘is a metaphor for’; visual metaphor; the figurative and the metaphorical (including here metaphor and synaesthesia). Before embarking on these seemingly disparate topics, I promised that they would play a part in a consideration of that other dimension of complexity in metaphor which, for the want of a better word, I described as richness. It is now time to make good on that promise.

My hesitation over the choice of a term to contrast with syntactic complexity arises from the fact that there are other words that could have been helpfully used. Metaphors can also be described as ‘vivid’ or ‘deep’, and these are perhaps more pointed than the rather vague ‘rich’. However, the agenda for the concluding section of this chapter has two items on it, and I needed a term which fitted them both. On the one hand, there is the issue of what makes particular metaphors powerful in some dimension or other. And, on the other hand, there are general issues about the status of metaphor itself. While it might have disadvantages, we can speak about both these issues in terms of richness: we can speak of what makes this or that metaphor richer than another; and we can speak about why (or whether) metaphor is an enrichment of language (and thought). What I propose is first to consider what counts as richness in metaphors, and then to turn attention more fully onto the question of the richness of metaphor. Along the way, I shall call on the various points made in sections 4.3–4.7.

4.8.1. The richness of metaphors

As already suggested, the simple subject-predicate metaphors so often used in philosophical discussions tend to be lower on any scale of richness, but it would be wrong to think this a result simply of syntactic poverty. Admittedly, ‘real’ metaphors, the kind that matter most to us, or at least strike us most forcibly in novels and poetry, and even scientific writing, are unlikely to be of the subject-predicate form. That is why it was crucial for me to show how the semantic descent account could take syntactic complexity in its stride. Still, even subject-predicate metaphors differ in their richness. Typical philosophical examples claiming human beings to be lions, foxes, or blocks of ice are impoverished—they tend to be moribund or dead metaphors—but, for all that it is familiar, Romeo's claim about Juliet is not like them. There is thus something we ought to be able to say about richness in metaphor which is independent of syntactical complexity, even if such complexity is often a contributing factor.

What I will say about richness will of necessity be sketchy. The subject is as vast as the study of literature itself, and I wouldn't even know where to begin that. However, adopting the perspective of the semantic descent account can shed some light on the kinds of consideration that might well figure in any such study, and I will content myself with some remarks about them.

Recall that in the discussion of dead metaphor (section 4.4), I used the resources of the semantic descent account to set up an ordering which ran from items we do not any longer or at all recognize as metaphorical, for example, ‘consider’ or ‘by and large’, to etymologically recognizable once-were-metaphors (‘reflect’), and idioms (‘kicked the bucket’), then to dead metaphors (‘lions’ and ‘foxes’), and finally to metaphors like (R) that are, by any standards, alive. What made this ordering possible is the fact that the semantic descent account requires live metaphors to call on both semantic descent and qualification for their intelligibility. In so far as one or both of these processes is absent, dormant, or works in a somewhat different way, we have ex-metaphor, idiom, and dead metaphor. But, as I noted at the time, the very fact that both processes are required by the category of live metaphor suggests that, by looking more closely at these processes, we can add a little order to that category.

As we have seen, some metaphors make very heavy demands on shared attitudes towards objects, attitudes which define what I called the ‘cultural’ significance of these objects. A second group, drawing on fairly straightforward recognition of what is common empirical knowledge, draw little from this kind of significance. Romeo's (R) is more of the former kind, but a more complex example might be: That time of year thou may'st in me behold
When yellow leaves, or none, or few do hang
Upon those boughs which shake against the cold,
Bare ruined choirs where late the sweet birds sang.¹⁶⁹

¹⁶⁹ White (1996: 154–8) makes much of Shakespeare's Sonnet 73 of which these are the first four lines. (See Shakespeare 1988: 760 for the text used.)

These lines could be used to illustrate both metaphor's syntactic complexity and the idea of extended metaphor, but I shall confine my remarks to the single metaphor that figures in the last line (though we need the first three lines to appreciate the subject or subjects of this metaphor). What I suggest is that even before we spend time on further elucidation and commentary, we have no trouble in recognizing that there is a kind of depth to this metaphor. Using the ruined choirs of a church to characterize trees in autumn (when these have already been used to characterize a stage in human life) calls on significances that perhaps could be articulated (by someone other than me) but which reach deeply into, dare I say, our collective psyche. One can see depth here, without necessarily seeing into it.

The following is an example of the second general sort of metaphor described above: Life is a dance floor, God is a DJ.¹⁷⁰ This metaphor requires (or presumes) fairly straightforward empirical knowledge: we need to know something about clubbing, and the loud, incessant music—controlled by the DJ—which dictates in some indirect sense the swirling movements of the clubbers on the dance floor. The metaphor here is certainly vivid, but it isn't deep.

It is usual to think that depth and vividness in metaphors go together, and it is certainly possible for a single metaphor to possess both features, but I am suggesting that, in calling on different resources, they are separable. Those we regard as vivid tend to make greater calls on the process of semantic descent; we are led by the words in the metaphor to some quite specific object, thereby leaving little room for doubt about its qualificational potential. In contrast, in those metaphors which tend to have depth, descent to an object, even descent controlled in certain ways by the words in the metaphor, still leaves a lot to be done. We know that Romeo uses that most prominent of objects in our sky to characterize Juliet, but there must be a sharing of the significance of that object before there can be anything like attunement in this characterization.

My talk of the different resources called upon in vivid and deep metaphors is liable to be misunderstood. I am not claiming that, when semantic descent bears the burden, we *thereby* have a vivid metaphor, or that depth *results* from reliance on cultural significances. It is perfectly possible to have limp metaphors which depend on semantic descent and shallow ones in which cultural background is fully engaged. All that I am claiming is that it makes sense to think of some metaphors as vivid, but not deep and that, in those cases where this does make sense, this can be explained by the relative roles of semantic descent and qualification.

The distinction between vividness and depth might be useful or interesting in discussions of specific metaphors, but it is not exactly the kind of thing that philosophers have at the forefront of their minds when asking after richness in metaphor.

¹⁷⁰ I saw this flyposted on a telephone switch box when I was walking in London. It is a quotation from a pop song.

What would be more like it would be considerations showing that a vivid and/or deep metaphor does something that simply cannot be done with plain language. Tempted though I am to say there can be no such *general* considerations, I do realize that this is bound to be disappointing. What is wanted is a way of describing the distinctive contributions of suitably rich metaphors, and a necessary precondition of describing such contributions is an answer to the question: is metaphor itself in some way special? Not wanting to disappoint, I turn now to this issue.

4.8.2. The richness of metaphor

How special or important is metaphor? This question is prominent in most philosophical discussions of metaphor, and it tends to be a divisive one. Some claim that metaphor is at best ornament, and at worst a devious rhetorical device apt to mislead. Most famous here is Hobbes's advice (itself given in terms of a metaphor, even if a familiar one): 'But for Metaphors, they are in this case ["rigorous search for truth"] utterly excluded. For seeing they openly professe deceit: to admit them into Councell, or Reasoning, were manifest folly' (cited in Blackburn 1984: 172). Others insist that, in some unspecified and mysterious way, metaphor is foundational to language and thought. Nietzsche is often cited, and the following is a particularly firm statement of his view: Truths are illusions which we have forgotten are illusions; they are metaphors that have become worn out and have been drained of sensuous force, coins which have lost their embossings, and are now considered as metaphor and no longer as coins. (From 'Über Wahrheit und Lüge im aussermoralischen Sinn', as cited in Cooper 1986: 258 and Blackburn 1984: 181)¹⁷¹

There are serious problems evident in both of these citations. Hobbes accuses metaphor of professing deceit, and he presumably has in mind here that thought, shared by all too many philosophers, that metaphors typically say the false: Juliet is most certainly not the sun; Tolstoy is long since grown out of infancy. In picking and choosing examples, one can make it seem as if metaphors are literal falsehoods, but a more cautious view is that, taken literally, they do not so much as possess sense. And, even if there weren't many other reasons to think so, this in itself should suggest that metaphors do not profess deceit, and therefore that they must be understood in some fundamentally different way.

The passage from Nietzsche is typically opaque, but philosophers tend to understand him as saying that ordinary assertions—'Truths', as he has it—involve dead metaphors. However, if this is all that Nietzsche is saying, it is not a very profound defence of metaphor. As noted earlier, those who find the claims made on behalf of metaphor overblown can perfectly well concede that metaphorical processes have left some mark on expressions we currently use. But this is a far cry from saying that metaphor itself is in any sense primary. Cooper 1986: 258 does cite a source (Bolinger 1980: 143–5) who claims that 'every thought and utterance we make is

¹⁷¹ It is no accident that I have found it convenient to borrow both the Hobbes and Nietzsche citations from Blackburn. He himself is following tradition in pitting Hobbes against Nietzsche when it comes to metaphor. I shall come to Blackburn's own view about whose side to take in this dispute.

charged with metaphor', but Cooper goes on to note: 'Since there clearly are literal utterances which we can distinguish from metaphorical ones, Bolinger's claim is absurd' (Cooper 1986: 259). So, Hobbes's unfairness seems to undermine his claim, and Nietzsche perhaps says something obviously true which can be over-inflated. Still, it is not easy to avoid feeling that these writers (and others) are taking sides in what is a substantial debate about the importance of metaphor.

What I think true, and will argue in the remainder of this section, is there is something right in *both* the Hobbesian and Nietzschean attitudes towards metaphor. Without going along with the letter of what Hobbes says, I think it true that metaphor is often mainly a stylistic ornament. And, without basing my claim on the corpses of dead metaphors which Nietzsche finds in our assertions—corpses scattered widely in dictionaries—I think there is good reason to think that metaphor has a kind of primacy that Cooper thinks absurd. However, I do not think much follows about the importance of metaphor from either of the claims I shall defend. I do realize that this sounds paradoxical, but I hope you will come to see that it isn't.

As you would expect, I think the key to getting all these matters straight lies in the semantic descent account. According to that account, metaphor is fundamentally a linguistic phenomenon. We render metaphors intelligible when we find certain words in them employed to access objects, and the objects thus made salient functioning as predicates, or as ingredients in predicates. Looked at from the middle distance, there isn't much difference between making an assertion using, as it were, only words, and making an assertion that is a metaphor. Nor is this only a matter of communication to others: formulating thoughts is no less a matter of choice between words and the processes of metaphor. Wanting to characterize Juliet, Romeo can reach either for words that are ready-made for this purpose, as when he says: (18) Juliet is a woman, or he can use words, not themselves functioning as characterizers of Juliet, but as devices which enable him to reach for an object which then takes on the characterizing role, as when he says: (R) Juliet is the sun.

This is not to deny that there are real differences between the literal and the metaphorical; indeed a large part of this book has been devoted to explaining them. It is rather to stress something that has also been central to my account: the continuity between the literal and the metaphorical. And the reason for stressing this continuity here is because of its consequences for the question of metaphor's importance.

Imagine, for example, someone asking after the importance or value of our ability, shown by (18), to use words. We might try: being able to use language makes possible expressing, communicating, entertaining, or recording our beliefs about how things are. But this is so blindingly obvious that one feels there must be more to the question than that. There are of course contexts in which the question, and these sorts of answer, have some point. Wondering about the evolutionary advantage that language confers on its users might be one such context. Yet even here language's

importance is so obvious that spelling it out, as above, would be more a reminder of what language is, than a serious answer to a serious question. Yet, when it comes to metaphor, a similar question is often assumed to be not only serious, but one actually requiring an answer. How come?

There isn't too much of a puzzle here. If you begin by assuming that metaphor is discontinuous with ordinary uses of language, even perhaps at odds with it, then it would indeed make sense to wonder why we go in for it. One is spoilt for choice here, but, given that he makes the issue of metaphor's importance his main contribution to the subject, I will use Blackburn's (1984) discussion as an illustration.

Beginning with what he calls the 'prosaic end of things', Blackburn concedes that we do use dead, or nearly dead, metaphors to express beliefs. But he insists the metaphor, in these cases, is just an indirect way of saying something that we could put directly, and without metaphor. To take his example: 'Bert is a real gorilla' yields that Bert is strong, rough and fierce. It does this because the word 'gorilla' has its normal meaning, and because people (wrongly) associate these features with gorillas. Any hearer aware of this can follow the metaphor to its intended interpretation. (Blackburn 1984: 172)

Given the eliminability that he finds in the use of dead metaphors, Blackburn can find no special ground for importance in them—unsurprisingly given what I said above—so he moves on to examples of live metaphor. However, a live metaphor is, by his lights: 'an invitation to explore comparisons. ... it is not associated with any belief or intention, let alone any set of rules, determining when the exploration is finished' (Blackburn 1984: 174). It is the sentiment expressed here that breathes life into the question of metaphor's importance. Obviously daft as it is to ask, just like that, whether being able to express or entertain beliefs is important, it is perhaps not so daft to ask about the importance of an activity that involves only what looks like the expression of beliefs. Put like this, it is unsurprising that the only way Blackburn finds to answer such a question affirmatively is by claiming: 'the value of a metaphor is essentially that of a means to an end. The end product is the appreciation of a literal truth or several literal truths. The metaphor suggests how to go about finding some such truth or truths' (Blackburn 1984: 175). And, though he does at one point imagine someone desperately trying to do better, someone who insists that 'there is a distinct, intrinsically metaphorical, way of understanding', he slaps this down by insisting: 'The man says the words [of a live metaphor] supposing that they mean something true, and that he knows what it is. More thought would disillusion him. So he simply believes he has a belief, when in fact he has none. He thinks he is rich when he has a cheque which he cannot actually cash' (Blackburn 1984: 176).

I would hope at this point that you would agree with me that Blackburn's conception of metaphor is hopelessly wrong, in respect of both dead and live metaphors. But the point here is not to argue with him, but merely to show how the whole question of whether metaphor is important depends on some such view as his. If you begin by assuming that the author of a metaphor uses words in what looks like an assertoric

manner, but that no assertion (except a bizarrely false one) is made, no belief is expressed, then you are bound to wonder why anyone would do that. Or why anyone would even entertain a metaphor as part of understanding how things are. However, if, instead, you have been persuaded by my conception of metaphor as special but continuous with other uses of language, then you should find it only puzzling that someone might ask whether metaphor is in general important. Sometimes it is, sometimes not; everything depends on the particular metaphor.

Alright, you say, when one stresses the continuity of metaphor with uses of language more conventionally thought of as linguistic, the question of importance can seem foolish. But, without denying the continuity, it is still part of the semantic descent account that metaphors involve a predicative use of *objects*. So, perhaps the right question about metaphor is this: is there a special value in using objects in this way?

This sounds like a better question, but I don't think it is. For the only answer that makes any sense is the same as the one I gave above: it depends. There are metaphors which, in terms of value, are pretty well indistinguishable from literal utterances. My point is got across by saying of some acquaintance that he has the personality of a traffic cone. I could have said instead that he is personality-free, or dull, but I chose to use a metaphor.¹⁷² Why? Well, it's funnier (even if a bit cruel). But we would be wrong to expect, on the basis of this example, some general answer that began: by using an object instead of a word, a metaphor can achieve ...

In other cases, the use of objects can appear to be special. Thus, we would struggle (in vain) to mimic the communicative or cognitive effect achieved by Updike's: The millennial fact for Hope is a large blank door that has slammed against her life, using only literal utterances. But I do not think that this shows that objects by themselves get up to tricks beyond the literal. It is certainly true that a deep and vivid metaphor can communicate a whole range of things all at once, and that we cannot replace it by any single literal utterance. These features are often cited as grounds for metaphor's attractiveness. But either or both features can also be true of literal utterances: in philosophy, but not only in philosophy, one finds literal utterances which communicate a great deal all at once, and which cannot be replaced by different formulations. The penchant of philosophers to begin, end, or merely

¹⁷² Blackburn thinks that dead metaphors are eliminable, and that is why they are so easily absorbed into our ordinary assertoric practices. But the metaphor here is not dead. Still, does the fact that we could have achieved much the same result using only words (as it were) show that the metaphor is eliminable? If by eliminable is meant 'the other expression has the same meaning as', the answer is most certainly 'no'. (I have also argued this is just as true of dead metaphors—the gorilla is not eliminable in this sense from the earlier claim about Bert.) If, instead, eliminable simply means that we could achieve the same general aim by using only words, then of course the traffic cone is eliminable. But the argument could as easily be run in reverse to show that speaking about someone as 'personality-free' is eliminable in favour of the metaphorical. The simple issue about this case is not whether the metaphor is or isn't eliminable, but that it is very ordinary, even if it is in some sense ornamental. Adapting Blackburn's own metaphor, it is a cheque, which we can use *instead of money* to pay bills, but it is not worth all that much.

support a complex argument by citing some philosophical authority verbatim is evidence of this.

So, in this first round, I find myself sympathetic with *both* the view that metaphor is sometimes largely ornamental, *and* with the view that it is sometimes irreplaceable. As I have insisted, it all depends on the metaphor. Still, there remains something to be said about the claim of metaphor to be fundamental, something that might justify us in thinking—what Cooper found absurd—that all our utterances are charged with metaphor.

As I have so often stressed, metaphor depends on both semantic descent and qualification, and this is true whether the metaphor is dead, obvious, vivid, or richly rewarding. In the case of dead metaphors, such as the one Blackburn cited earlier involving the gorilla, we may not actually need these processes to confer intelligibility; but their background presence is shown by the ease with which even the most dead of metaphors can be ‘revived’ in the right context. It may even be true (or so I argued) that descent is operative, but not qualification in the full sense. One way to understand what is going on in ‘Bert is a real gorilla’ is that there is descent to the gorilla, but this object is then used to refer metonymically to the property of being strong–rough–fierce. Qualification thus is not needed for our understanding what the complete sentence says about Bert; it only figures in the metonymic move from the gorilla to a property. Earlier in the chapter, I speculated that this is a better way to understand some cases and, in respect of ‘gorilla’, it would avoid having to attribute false beliefs to users of this dead metaphor as Blackburn does. When the metaphor is live, whether it is obvious, vivid, or rich, both processes are engaged, but when dead, qualification can be sort of short-circuited by descent.

All that said, I also considered the very special case in which we have qualification, but not descent. From the very beginning, I stressed that qualification without descent is not metaphor. Seeing a large tree felled by a strong wind as a source of information about how to conduct one's business affairs, or seeing bizarrely constructed paling as a source of information about one's poetic efforts, are clear cases of the phenomenon of qualification, but they are not metaphors. Nonetheless, given its crucial importance *to* metaphor, it is intelligible why someone would think that qualification was enough *for* metaphor. (The very fact that we can speak of one object being a metaphor for another can confuse the issue, though, as I indicated earlier in the chapter, descent is operative here; we are, after all, *speaking* of these objects.)

Against this background, recall my scattered speculations about the role of qualification in the development of categorization. I emphasize the word ‘speculations’, but the point now is not to elaborate on, or defend, them, but rather to explain how they might fit in with the idea that metaphor—or rather a central process in metaphor—could be fundamental to language (and thought). If it were true, for example, that words like ‘ewe’ are the special result, not of our noticing similarities, but of our penchant for taking objects as qualifiers of others, then it would offer a reason, albeit slightly confused, for thinking that metaphor was at the heart of our linguistic practices. This would not be because ‘ewe’ was a dead metaphor, since

strictly speaking it isn't a metaphor at all. So, this would rule out the letter of Nietzsche's words. But if it were true that 'ewe' is a stand-in, a word-object which, as a result of a process that began with the qualification of one object by another, has come to refer to the property of *being a ewe*, then perhaps the spirit of Nietzsche's words can be preserved. We would have found a way to make good on the initially absurd-sounding suggestion that language and thought is charged with metaphor. Not every part of language, nor all thought: that would indeed be absurd. Nor, as I have been careful to say, would it really be *metaphor* that was strictly responsible for this charge. But, even if we were to say only that a good deal of language and thought can be understood by appealing to a central process in metaphor, perhaps that would be enough to satisfy Nietzschean supporters, while at the same time not treading too firmly on the toes of the Hobbesian detractors.

5

Competitors

5.1. Introduction

As I have insisted throughout, there is more right than wrong in many well-known accounts of metaphor. That said, I have not gone into great detail in expounding these views. This is partly because accounts like those of Black, Beardsley, Goodman, Searle, and Davidson are so familiar, and partly because of the minimalist strategy which governed operations up to Chapter 4. However, in this chapter, I will allow myself to indulge in some expository and textual detail, though it will not be targeted on these well-known accounts, but on three which may be less familiar. There are two reasons for this. First, there is the matter of testing the classificatory scheme given in Chapter 1. If that scheme holds up, these views, like those mentioned in that chapter, must fall under one of the three heads. (I shall argue, that they do.) Second, and much more importantly, in one way or another, these views circle around my own; in studying them, I have had the feeling that they were in places right on the edge of thoroughly anticipating what I have been thinking for many years. That these writers never did come up with my view of metaphor is what made this book possible, and, in a way, I am grateful for that. But, in wanting to support my account as strongly as I can, these other conceptions of metaphor are an invaluable resource. As I see it, these views reflect the phenomenon of metaphor well enough to do more than make it recognizable: they help us to see more deeply into its nature than has been possible before. Yet the reflection in each case is distorted, not dramatically, but enough to blur certain features. If I can show that my own view acts as a kind of correction to these other views, a correction which, once made, gives us an undistorted view of those very same features, then I would expect this, added to the material of the preceding chapters, would count as further strong evidence in favour of the semantic descent account.

Though choosing from among the growing number of highly detailed accounts that have been published in the more recent literature is not easy, three stand out as especially relevant to my purposes. Two of these count as Content Sufficient, and one takes the Alternative Message route. That there is no representative in the third category shouldn't come as a surprise. Davidson is pretty much alone in propounding a No Content/No Message view, but that is not a mere contingency. Given that any view in this third category is defined by its opposition to the projects of those in the other two categories, we shouldn't expect more than one representative in this

category, though of course arguments for the view will differ from one writer to another, as will suggestions for the role of metaphors generally.¹⁷³

5.2. The Demonstrative Account

In his recent book, *Metaphor in Context*, Josef Stern has added a great many details to what is essentially the account he published in a series of articles centred around Stern 1985. In the book, there are a number of ramifications and even alterations to the original account, but, as the aim of my discussion is not point-by-point exegesis, I shall not pause to consider the ways in which the view has changed. In fact, for obvious enough reasons, everything that I shall say will be based upon the book, and even then it will be a bare-bones outline of the view.

Imagine that the familiar sentence: (R) Juliet is the sun, had been uttered in a completely different setting. Imagine that it was said by someone who prefaced the remark with a pointed reference to the myth of Icarus. Or imagine it said by someone just after a conversation about the dangers of looking directly at a total eclipse. Clearly, we have a sense that the utterance of (R) would in each of these cases be interpreted differently by any reasonable hearer, and the obvious fact of this kind of variation is the starting point for Stern's account. (Note that it could also be used as the starting point of many other accounts; Stern's view is distinguished from others, not by where it begins, but by its destination.) What Stern asks us to notice is that, while there seem to be indefinitely many interpretations of (R), there is a pattern to these interpretations which suggests something like a constraint at work. In his words: The moral I draw is this: There may be little systematic or predictable so long as we look just at the particular contents of the different metaphorical interpretations one by one, but at one level of abstraction more abstract ... metaphorical interpretation does seem to follow a pattern and to support predictions. Same expression, same context, same interpretation; same expression, different contexts, different interpretations. (Stern 2000: 12) Leaving aside the fact that metaphor is at issue, the kind of patterns that Stern describes are of course perfectly familiar: demonstratives and indexicals are words which have different content ('interpretations') in different contexts. Or, in a slightly more focused formulation, the items which they bring to the content—their referents—vary with context of utterance. It is this parallel which lies at the heart of Stern's account. Again, in his words: To return now to metaphor: instead of allowing its context dependence to be an obstacle to its semantic candidacy, I'll argue that the key to its satisfactory semantic analysis is to

¹⁷³ For example, Cooper 1986 shares Davidson's No Content/No Message view though he does not share all of Davidson's arguments for that view, nor does he think of the role of metaphor in exactly the way that Davidson does. From what I can tell, Blackburn 1984 is also a supporter.

embrace its context-dependence. To go one step further, I shall treat metaphor as a type of context-dependent expression of the same general kind as the demonstratives. What is needed for such an account is to show that metaphors and demonstratives share the same formal structure, to isolate the contextual parameter (like the speaker for 'I' and time for 'now') that determines a metaphorical interpretation, and to specify the rules that determine the contents of metaphors in each of their contexts of utterance. (Stern 2000: 14–15) I'll return to Stern's insistence that his account is 'semantical'. By and large that is not really all that important, except of course to classification of the view. (And even then it is not all that important.) But taking things one step at a time, I must outline what sounds a very bold proposal indeed. After all, given the trouble that metaphor has caused, especially for accounts that claim to be semantic (or 'Content Sufficient', to use my term), talk of 'rules that determine the contents of metaphors in each of their contexts of utterance' sounds almost too good to be true.

The background story about demonstratives and indexicals that Stern considers both correct and fundamental is Kaplan's. And it is not a short story, though here it will be abbreviated. Essential to Kaplan's account is the idea that demonstratives and indexicals have what he calls 'character', a notion that is somewhat off to the side of more traditional ideas of meaning or sense. Using 'I' as the model (which most people, including Stern, tend to do), the idea is simple enough. Imagine the sentence: (1) I am going to visit Hungary, as uttered (*sotto voce*) by a number of the tourists on a bus pulling out of Victoria Station, and bound for points south and east, including Budapest. Someone familiar with English understands (1) without difficulty, but that same person would have a great deal of difficulty saying whether (1) was true. No great mystery here: (1) is not a sentence that can be evaluated for truth just by understanding it; its truth depends crucially on the referent of 'I'. It is true only as uttered by each person on the bus who is actually going to Hungary. It is not true of someone who didn't utter it (or at least think it), and it is not true of anyone who may have uttered it, but who gets off at the wrong stop. As Kaplan sees it, sentences that have words like 'I' in them (and 'here', 'this', and others) force us to recognize a more kinky route in the move from understanding to truth conditions. The truth conditions of (1) (often called 'content' by those of the Kaplan persuasion) wait on our having the referent supplied. However, there is something that we, as speakers of English, grasp ahead of truth conditions—something which in fact constrains the truth-conditional content that we eventually arrive at. This extra element is the *character* of 'I', something that we understand in (1) even when we are still in the dark about the identities of the Hungary-bound tourists.¹⁷⁴

Character is a constraint on content, and for 'I', we can actually state the rule which realizes that constraint (this is a version that Stern gives): (I) Each occurrence of 'I' directly refers to the agent of the context, e.g. its speaker. (Stern 2000: 81)

¹⁷⁴ As will eventually be seen, Stern's view is Content Sufficient, but 'Content' here is not Kaplan's 'content'. I shall deal with this possible source of confusion below.

Stern's wish to see metaphor as context-dependent—as just like indexicals and demonstratives—suggests that each metaphorical expression acts like a constraint on the final content expressed by the sentence in which that expression occurs. Just as rule (I), in operation in a context, constrains our view of what (1) actually says—its truth-evaluable content—so, Stern suggests, for any metaphor, there can be a rule which, taken together with appropriate contexts (and knowledge of the non-metaphorical elements of the relevant sentence) will yield a truth-evaluable content.

As noted, this is a strong claim. Moreover, Stern recognizes that the appeal to rule (I) is not yet a sufficient basis for it. The reason is simple: any rule of metaphor would have to deal with an indefinitely large range of lexical items (not to mention constructions out of lexical items), whereas rule (I) is tailored to just one fixed item, the word 'I'. However, help is at hand in the form of Kaplan's work on demonstrative descriptions. The details here can be quite tricky—especially as they involve the interaction between the different kinds of demonstrative construction—but I don't think that my resolve to keep things simple will render the summary below inaccurate.

Take any definite description, for example, 'the house at the end of the lane'. Kaplan's suggestion is that this description, which as a singular term is certainly non-rigid, can be converted into a rigid, demonstrative-like, singular term. This conversion is effected by prefixing the description with a 'Dthat' operator. Thus, (2) Dthat ['the house at the end of the lane'] is for sale, is a sentence whose truth-conditional content is now quite different from its non-'Dthat' version, (3) The house at the end of the lane is for sale. In sentence (3), we use our knowledge of meaning to ascertain character—a constraint on truth-evaluable content—but this character is fixed for all the different contexts in which (3) is evaluated for truth. The fixedness of character here is the simple consequence of the fact that we understand the meaning of the words in the definite description in (3) without needing to appeal to anything contextual. What we have in (3) is a sentence whose character is not context-sensitive, even though the actual truth or falsity of this sentence most certainly is.

In contrast, 'Dthat' in (2) forces the interpreter to appeal to context, not merely as a part of truth assessment, but in order to determine the character of the sentence in the first place. Just as 'I', in meaning (approximately) 'whoever is the speaker in this context', makes a context-sensitive contribution to character, 'Dthat' forces us to think of the whole, namely, 'Dthat ["the house at the end of the lane"]' as making a context-sensitive contribution to the character of (2). Moreover, and this is the crucial bit for us, the context-sensitivity here is not merely that of certain lexically fixed items like 'I' and 'this', but is imposed on an endlessly differing set of lexical items, namely definite descriptions.

The context-sensitivity that ‘Dthat’ imposes on a definite description shows up, as it does for ‘I’ and ‘this’, in a rule (as before, I am using Stern’s version):(Dthat) For every context *c* and every definite description ϕ , an occurrence of ‘Dthat [ϕ]’ in *c* directly refers to the unique individual (if there is one) denoted by ϕ in the circumstance of *c*, and to no one otherwise. (Stern 2000: 100) Like the case of ‘I’, where what matters to the truth-evaluable content includes a referent determined in a context-sensitive way via the rule (I), so here what matters is the referent supplied by the context, via the rule (Dthat). In essence, the ‘Dthat’ operator turns an arbitrary definite description—an expression whose meaning or interpretation (i.e. contribution to truth-evaluable content) is determined non-contextually—into an expression whose contribution to content explicitly varies from context to context. As Stern 2000: 103 puts it: “‘Dthat’ provides us with a technique ... to import context-sensitivity into an interpretation where previously there was none’. And this is precisely the kind of thing he requires for his view of metaphor, namely, a character-creating operator that can take various linguistic expressions as arguments, and whose outputs are contextually differing contributions to content. There is, however, still one obstacle to moving too quickly from the rule (Dthat) to a rule for metaphor.

The argument expressions in ‘Dthat’-constructions are definite descriptions, and the rule (Dthat) uses this fact in its operation. In order to work out the contribution to content of, for example, ‘Dthat [‘the house at the end of the lane’]’, one must take into account the denotation of the ‘the house at the end of the lane’. That is, one must in this case work out what item (if any) this description applies to. This kind of move has no place in Stern’s account of metaphor. What he wants instead is a character-creating rule that works directly on metaphorical expressions, rather like the rule for ‘I’. Stern makes it absolutely clear that it is the *words*, the lexical items, on which his operator works. For example, he notes with approval just this feature of Black’s story about associated commonplaces when he writes: ‘He does not explain why he says it, but Black is also right that the relevant beliefs must be associated with the *word* ϕ and not just with its extension or intension’ (Stern 2000: 108; his emphasis). The importance of this concentration on metaphorical *expressions* will become apparent shortly.

What Stern then wants for metaphor is a rule that models itself on a combination of (I) and (Dthat). The (I) rule gives you a contribution to content that is context-sensitive, and it applies directly to expressions like ‘I’ and ‘this’. But the ‘I’ rule only governs a single fixed lexical item, and this won’t do for metaphor. The (Dthat) rule, however, shows how one can graft a special kind of context-sensitivity onto a range of different items which, to begin with, do not possess such context-sensitivity. However, unlike the rule for ‘I’, the rule for ‘Dthat’ does not apply simply to expressions (i.e. definite descriptions), but rather to the referential function of these expressions, and while this too won’t do for metaphor, Stern’s idea is that one can imagine a rule with the useful features of (I)-type and (Dthat)-type rules combined.

The transition that Stern makes from demonstrative and descriptive context-sensitivity to metaphor is far from simple, and it depends on various further notions. However, given the level of detail appropriate here, one can bypass a lot of the discussion by going straight to the rule that Stern proposes for metaphor, afterwards working out how it applies to an example. The rule is: (Mthat) For every context c and for every expression ϕ , an occurrence of 'Mthat [ϕ]' in a sentence $S \dots$ in c (directly) expresses a set of properties P presupposed to be m -associated with ϕ in c such that the proposition, $\langle \dots P \dots \rangle$, is either true or false in the circumstance c . (Stern 2000: 115) The rule (Mthat) is intended to show *both* how any expression can be turned from its ordinary self into an expression which is context-sensitive *and* that when this context-sensitivity is filled out, we get precisely what one would expect from that expression used in a metaphor. In effect, the 'Mthat' operator imposes a context-sensitive *structure* or *form* on an expression, and the rule governing the operator points to the way to filling in this structure so that, in a relevant context, it possesses an appropriate metaphorical *meaning or content*. (Labels are a bit of a problem here: strictly speaking, what one ends up with is a determinate Kaplanesque character which, in a certain circumstance, will impose this or that Kaplanesque propositional content. But the basic idea is clearer without this adherence to the special terminology of Kaplan's views, and Stern himself is no stickler for these labels.)

An example illustrating the (Mthat) rule will certainly make it clearer. Take the expression 'is the sun' as it occurs in: (R) Juliet is the sun. When the rule (Mthat) is applied to it, the familiar context-invariant meaning drops away, and in its place we have a structure whose meaning varies with context. How do we know what this meaning is? This is the longer bit of the story that is condensed into the (Mthat) rule.

According to Stern, 'is the sun' has associated with it a set of properties which I shall call N . These will include such things as: is at a great distance, is a source of warmth, makes one feel good, is what we want to see on holiday, is a source of light, is hot, burns, is gaseous, is yellow, rises and sets, is at the centre of the solar system, is dependable, etc. Stern aims to be quite liberal about the kinds of presupposition that are the sources of the properties making up N ; they can be stereotypical beliefs, ideas about typical or normal uses of the expression, or even ideas about recent or prominent special uses. However, in the Shakespearean context of the utterance of (R), certain of these properties will be inappropriate to the content that Romeo intended to convey. Indeed, Stern assumes that only a subset of N will capture that content, and this is the set he calls P . Why then is there no mention of N in the rule (Mthat)? Because, given the fact that P will be selected in a context, Stern insists that there is no harm in imagining that, from the very beginning, $N = P$. To be sure, without any specific context in mind, N is likely to contain all sorts of properties that

are irrelevant to its use in the context of (R). But we can simply assume that any set of properties that is of interest to us will be drawn up context-by-context, and this turns the selection function into one of identity.

Here I have been careful to explain N separately from P because it avoids certain problems. Treating N as I have, one can see that m -associated properties are quite various and are not essentially tied to any particular metaphorical use, or even to metaphor. Indeed, Stern would have been better advised to call them simply 'associated properties', since the prefix 'm-' (i.e. 'metaphorical') forces him to dodge complaints about the circularity of his account.

Now, given that P is the right set to associate with 'is the sun' in the context of (R), we can take 'is the sun' to attribute just these properties to Juliet. Yet in contexts different from the Shakespearean one, some other set, say P^* , might serve instead. As I have stressed, Stern's key claim is that metaphors are expressions which in different contexts make quite different contributions to the truth-evaluable (propositional) contents of sentences in which they occur. Hence, though the sentence: (A) Achilles is the sun, calls on the same expression 'is the sun' that occurs in (R), its contribution to the truth-conditional or propositional content of (A) will be different from that of 'is the sun' in (R). This is because, in the appropriate context, the properties associated with 'is the sun' in (A) will not be the same as those associated with this expression in (R). (Stern presumes that in the context of (A) the associated properties will include items along dimensions of power and intensity which would be inappropriate for 'is the sun' as applied to Juliet.)

5.2.1. Problems

The bare-bones presentation of Stern's view does not even begin to do justice to the rich range of consequences and applications which fills out his view. However, while a great deal of what Stern says about metaphor and metaphor phenomena can be discussed and evaluated independently of the demonstrative/ contextual kernel of the account, the kernel is crucial. And even my bald presentation of that kernel puts us in a position to ask this crucial question: does the central demonstrative/contextual idea bear the weight which the rest of his discussion imposes on it?

There are substantial obstacles to giving an affirmative answer to this question. First, the models in Kaplan's work that Stern calls on do not transpose easily, if at all, in the way that Stern requires. Indeed, one may well wonder whether the rule Stern gives us for 'Mthat' is even coherent, given its origins in Kaplan's rule for 'Dthat'. Second, there are problems about the m -associated properties which, in the (Mthat) rule, give substance to the demonstrative contextual account. These problems concern both the plausibility of the way in which they are selected in any given context, and more fundamentally, their plausibility in relation to the phenomena of metaphor.

Most of what I want to say about Stern's proposal revolves around the second of the above points. But a few remarks about the first will prove useful.

5.2.2. Kaplan's rules and metaphor

As noted earlier, Kaplan's (I) rule, as presented by Stern, describes the context-dependent character of a particular lexical item, the expression 'I'. Informally, what rule (I) displays is the contribution that 'I' makes to the propositions expressed by sentences in which it occurs. What the (I)-rule makes plain is how the contribution of 'I' differs from the contribution of, say, a word like 'horse'. The latter's contribution to propositions expressed by sentences in which it occurs is both determinate, and not dependent on context, whereas the contribution of 'I' is determinate and context-dependent at the same time. This point is simple enough even if the terminology can get out of hand. When we read: (4) I am tired, we know in some sense what is being asserted, even though we might not know which truth, or whether a truth, is expressed, simply because we don't know who produced (4). But if we read: (5) A horse is an animal, we know in the appropriate sense what is being asserted, and we also know which truth it expresses. Context doesn't come into it, at least not in the way it does for (4).

However, the (I) rule is not a good model for bringing about context-sensitivity in metaphor: accepting it as such would require a separate rule for the indefinitely large number of expressions that might pop up in metaphors, and even this wouldn't be enough. For, unlike 'I', there would be no natural context-sensitivity in those expressions; 'horse' can be used metaphorically in various contexts, but there is no actual context-sensitive hook in 'horse' to serve as a point of attachment for these metaphorical meanings.

These considerations prompt Stern to look to Kaplan's discussion of definite descriptions, and the 'Dthat' operator. As noted earlier, this operator has the effect of inducing context-sensitivity in expressions which, on their surface, lack it. Moreover, 'Dthat' manages this in respect of any of the indefinitely large number of descriptions that we might come up with; it is not tied to a single lexical item. So, can we appeal to 'Dthat' and the (Dthat)-rule as a good model for Stern's 'Mthat' as embodied in the (Mthat)-rule? Sadly, the answer is 'no'.

It is absolutely essential to (Dthat) that the expressions which it governs are of a particular kind; one would not be wrong to say that they are of a particular semantic kind. Definite descriptions can be as various as our imagination allows, but they are all *definite descriptions*; that is, they all make the same *kind* of contribution to the truth conditions of sentences in which they occur.

Don't expect me to spell out here what that contribution is. There are endless debates about these phrases, and Kaplan's is only one contribution to those debates. All that I am saying, or need to say to make my point, is that, whatever view you adopt, it will give these items one (or perhaps a couple) of semantic functions. Kaplan's (Dthat)-rule depends on this. What this rule does is to turn any definite description into a device whose contribution to truth-conditional content is context-sensitive (like 'I'), but whose meaning (contribution to 'what is in some sense said')

remains as variable as the description on which it operates. That it manages this trick is due wholly to the predictable nature of definite descriptions as expressions which, say, can be used to refer to, or call attention to, those things which uniquely satisfy the predicates they contain.¹⁷⁵ Kaplan's 'Dthat' cannot be attached to just any expression; not only would the rule fail to induce context-sensitivity in, say, a prepositional phrase, it would make no sense to apply it to such a phrase. Yet it is explicit in the (Mthat)-rule, and crucial to Stern's aspirations for it, that his rule takes as its field of operations 'expressions'—*any* expressions. The 'Mthat' operator is somehow supposed to turn expressions of any syntactic or functional type into items which have a special sort of context-sensitivity. The idea that we could have an operator which could be intelligibly prefixed to expressions of any type is not absurd: think of quotation marks which can be prefixed (and suffixed) to any expression.¹⁷⁶ But the (Mthat)-rule is clearly not like this, for the rule explicitly speaks about a context-sensitive link to *properties*, and properties are surely the kinds of thing we take to be expressed by predicate expressions. Unless all metaphor expressions are predicates, the rule (Mthat) is somewhat confused. For, as my examples in Chapter 4 made abundantly clear, metaphors can take on virtually any syntactic shape, and it is odd to speak of prepositional phrases, transitive verbs, and whole sentences, among other cases, as expressing properties—as functioning, in other words, the way that a predicate expression like 'is a horse' does. So, however suggestive the parallel might be with the rules that Kaplan's uses to deal with demonstratives and descriptive phrases, Stern's (Mthat)-rule is deeply flawed. It suffers from a kind of incoherence: the rule is simply not well-defined for the range of things it purports to take on.

5.2.3. Associating properties

It is certainly worrying that the central idea of *Metaphor in Context*—the interpretative (Mthat)-rule for metaphor—is not strictly coherent. But it might be thought that some kind of technical fix could remove or lessen the problem. Unfortunately, there is a second area of difficulty, connected to the first, which is far from being of merely technical concern. Here I refer to the idea that (Mthat) demands of a metaphor expression that it can have a set of 'm-associated properties', differently selected in different contexts.

There are three more specific problems here. First, it is far from clear what it means to 'associate' a property with an expression. On the one hand, it is unclear why it is properties, not objects, facts, events, or any other sort of thing which figure in these associations. On the other hand, it would seem that there are endless associations that a word or phrase (an expression) might have which have nothing at

¹⁷⁵ I trust that everything just said is intelligible, even without my having taken care over Kaplanesque labels like 'character' and 'content'. Also, remember that, though it might sound as if I am committed to this or that view of the function of definite descriptions, this apparent commitment is only of the 'let us suppose' variety.

¹⁷⁶ In a discussion circling around some of the issues raised by 'Dthat', Stern (2000: 103) does mention quotation. But he never links it to his own discussion of 'Mthat'. As you might imagine, the fact that quotation—something very much related to semantic descent—pops up here is no coincidence. See later in the text for more on this.

all to do with metaphor. For example, what if one associates the property of being my granny's favourite way of putting it with some expression?

Second, even assuming that Stern's view is restricted to predicate expressions, and not just expressions in general, and that he doesn't include such things as 'favourite expression of my granny' as live property candidates, there is still this problem: the properties one associates with predicates are unlikely to include those appropriate to the use of that predicate in a metaphor. For example, even assuming it is plausible to think of 'is the sun' as associated with, say, the property of being a source of warmth, it is far from clear that this is something that could be truly said of Juliet. As I have mentioned more than once, the sun's warmth-giving—that very property—is not one Romeo is attributing to Juliet in (R), even if he is there attributing properties to her, and even if one of these is that she is a warm person. It takes effort to understand the transformation (if that is what it is) of the sun's warmth into emotional warmth, effort which is precisely what an account of metaphor should explain.

This last point leads immediately to the third problem. Stern gives lots of examples of metaphors, but they are typically cases in which Stern and his readers would be able to make reasonable guesses about relevant metaphorical properties. So, when he points out that this or that property figures in the m-associated set for a given context, he is simply working backwards from his own understanding; he is inserting into the set of properties associated with 'is the sun', the property of giving emotional warmth. But, though it would be unfair to require an account to be wholly predictive of metaphor effects, there does seem something wrong with its being wholly ad hoc. It would have been nice to be given some independent idea of why, in various contexts, we can expect certain metaphor-relevant properties to figure in the m-associated set.

That these difficulties arise for Stern's account shouldn't be surprising, given that it has pretty much the same agenda as accounts like Black's. To be sure, Stern models his metaphor-inducing operator 'Mthat' on Kaplan's work on demonstratives, and there is nothing like this in Black, Goodman, or the other similar Content Sufficient theorists. Yet, as I have argued, the appeal to any such model is fraught with difficulty. And leaving the model on one side what we find is this: 'Mthat' is essentially an *explicit* instruction to create metaphorical meanings, in relevant contexts, for expressions to which it is attached. In other Content Sufficient accounts what is recognizably this same instruction is only implicit. This is a real difference, but not a crucial one, and other similarities are manifest: Stern and the Black-style Content Sufficient theorist insist that it is *expressions* which are the focus of our efforts in regard to metaphor, and both see the goal as that of arriving at some set of properties associated with expressions. Moreover, both regard these properties as replacing those associated with the literal meanings of the expressions with which one starts.

Given this background, we can understand the basic shape of the difficulties one finds with Stern's view. One has the suspicion that, for all the talk of 'expressions', it is only really predicate expressions that are covered by the account. This is a more

obvious problem for Black, and it is hidden somewhat in Stern's story by the apparently general nature of the rule governing 'Mthat'. But, given that this rule speaks of associated *properties*, it is clear enough that it was formulated with predicate expressions in mind. Following hard on the heels of this are the problems of telling an even minimally explanatory story about the linkage between predicate expressions and the properties that figure in the contextually created, metaphorical meanings. Black speaks about 'associated commonplaces', Stern about 'm-associated properties', Goodman about 'schema transfer', but in the discussion of actual examples, more or less suitable properties tend to appear miraculously in the relevant lists, even if one might doubt that they are commonplaces or associated in any natural way with the predicate *expression*.

Stern does at one point consider a slightly different, more promising, route from expression to m-associated properties. For this, he drafts in Goodman's notion of exemplification, and suggests that many metaphors, including Romeo's, are what he calls 'metaphors of exemplification'. He writes: When Romeo utters [R] Juliet is the sun I have argued, first, that its content is not that Juliet and the sun resemble each other (in some or another respect) and, second, that it is indeed unnecessary to detour through an underlying judgement of similarity to determine its content. Instead Romeo is saying that Juliet has a certain property—call it *P*—m-associated with the predicate 'is the sun' and that the crucial explanatory factor that generates *P* is its salience in the context. I want now to propose that, for at least one class of metaphors (like [R]), the salience of *P* derives from the fact that it is presupposed to be exemplified by the object(s) that is (are) the extension of the literal vehicle of the metaphor. (Stern 2000: 156–7) Leaving aside the details about presupposition and salience, the crucial points here are these: 'is the sun' is described as having an extension, namely, the sun; this object exemplifies a certain property *P*; and it is *P* which we understand as applying to Juliet. Why isn't this just as problematic as any other case? That is, given that the sun would normally be taken to exemplify various astronomically related properties, how come it figures in a predication about Juliet? Continuing the above, Stern writes: Which properties are exemplified by the referent of the literal vehicle of the metaphor? What any object exemplifies, we said, is relative to the other elements in the exemplification schema to which it belongs and to the field of features sorted by the schema as a whole. Metaphors of exemplification are no different; we therefore need to identify the schema and field of [R]. Let's look again at the context of [R] in Shakespeare's play. Here follows the full text surrounding Romeo's remark, a text in which lots of play is made with the moon and sun in the context of various human projects and emotions; the moon is envious of the sun, already sick and pale with grief, and would kill it, etc.

After discussing this passage at length, Stern concludes: 'What "the sun" expresses metaphorically depends on what the sun, relative to the scheme set up in the passage, exemplifies.'

Does this story help? Only a little. As I have argued, Goodman's notion of exemplification is a species of the genus *referential*, but it is not merely a referential relation. That a predicate can refer to an object which can then, in turn, refer back to the predicate (or property) is certainly a genuine linkage between a predicate expression and a property. But this linkage is only of some use if the backwards reference takes us to some property or properties other than that which the predicate would usually (literally) be taken to express. Thus, 'is a shark' expresses a certain property (*being a shark*), and a shark (the object) could be said to exemplify that property. But if 'is a shark' is used in a metaphor, we need the resultant property to be one that is different from that expressed by 'is a shark'. Admittedly, if the discourse context of a metaphor invites us to think that some property other than being a shark is at issue, we can extend our search. And it is not unreasonable to suppose that some property of a shark, other than simply its being one, might also be exemplified by a shark, for example, the property of being a voracious indiscriminate eater or a stealthy hunter. But we are still going to need some story about how these latter properties help us understand: (6) Harry is a shark. Obviously enough, this sentence suggests that Harry is someone who strenuously pursues his own interests without regard to others, and who cannot be trusted. It is after all a cliché that some such thing is being said about Harry by (6). But sharks do not exemplify these properties. The properties they do exemplify are in some sense close to them: we could even say that a shark pursues its own interests without regard to others, and that a shark is not trustworthy. However, in so saying, we would be speaking just as metaphorically about the animal as we would be in using (6) about Harry. For while the properties at issue are close, they are not close enough. Or, rather, they only become close enough once we understand them to be metaphorically linked, and this makes it hopeless to use their linkage to explain what is going on in metaphor.

It is unsurprising that there is a difficulty here because it is in effect the same as the difficulty that infects numerous appeals to similarity made in the attempt to account for metaphor. However tempting it is to say that (6) describes Harry as similar to a shark because of certain shared, contextually salient properties, a clear-headed look at the properties of sharks and the properties of Harry will show that any such sharing is a consequence, not an explanation, of the metaphor in (6). That this issue should surface in connection with Stern's view is also unsurprising. Though he eschews the direct appeal to similarity one finds in the simile view, he does give similarity a large role in generating the set of associated properties that figure in the (Mthat)-rule. He writes: 'I shall concentrate on one class of metaphors based on exemplification, a symbolic relation that I'll argue is a source of the ... similarity judgements that are said to ground metaphors' (Stern 2000: 146).

And he then embarks on a defence of similarity, not as capturing what is asserted by a metaphor, but as aiding us in the generation of sets of contextually sensitive m-associated properties.

5.2.4. The demonstrative account and semantic descent

I have gone on at some length about Stern's account in order to make it easier to understand its startling affinities—and disaffinities—with my own view.

When I first came across Stern's view in his 1985 article, my thought was that it had wholly anticipated what I had up to then only presented in informal talks about metaphor.¹⁷⁷ His 'Mthat' operator seemed merely a notational variant of my down-arrow notation, and I expected his talk of demonstratives to be in some way a version of semantic descent. However, I had skimmed the article too quickly and a closer look revealed that I was of course wrong on both counts. But there was still something right about my sense of what Stern was aiming for. Moreover these aims are ones that I share.

To start with, given my resistance to Alternative Message and No Message views, and my insistence that the original scheme of classification is exhaustive, my account must be Content Sufficient. Nor have I tried to hide this fact. Like Stern's version of Content Sufficiency, the semantic descent account aims to provide, context-by-context, alternative meanings linked to relevant expressions in metaphors. But radically unlike Stern's account—and those of Black, Beardsley, Goodman, and many others—my account does not seek to carve out new meanings *for those expressions*. For example, where Stern would have us set off on a search for some different set of properties associated with 'shark', that is, a new meaning for 'shark', in order to make sense of (6), my account rests content with the ordinary meaning of this expression. In my account, 'shark' rather satisfyingly means shark, and it is not the word but the fish itself which is the focus of our efforts to make sense of (6).

There is also a similarity between Stern's 'Mthat ...' and my '↓ ... ↓'. Each can be seen as requiring some kind of departure from the usual functions of the expressions to which they apply. But this similarity is superficial. Stern's operator applies to a word or phrase in a metaphor, and its purpose is to show how to generate, on a context-by-context basis, a metaphor-friendly replacement meaning for this word or phrase. In contrast, the down-arrow device marks a descent from the linguistic level—from expressions—to the objects which these expressions are about. When this is taken together with my account of the predicative function of objects, we do indeed get what could be thought of as alternative 'meanings' that help us to make sense of metaphors. But, unlike the proposal in Stern's and all other Content Sufficient accounts that I have come across, these are not replacements for the meanings of the original expressions.

As already noted, Stern gives a prominent place to what he calls metaphors of exemplification. These mandate a central, intelligibility-conferring role on the

¹⁷⁷ I first presented the account in talks at St Andrew's University and at Stirling in 1984. Though rudimentary, and not described as semantic descent, the basic ideas of the account were present in those talks, and had already figured in two years of desultory thinking about metaphor.

things which are referred to by some metaphor expressions. Thus, for Stern, the sun, and not merely ‘the sun’ figures crucially in helping us with Romeo's (R). However, while this makes Stern's account appear tantalizingly close to mine, this appearance too is deceptive. In Stern's explanation of metaphors of exemplification, objects like the sun are pressed into service in generating replacement meanings for the original expressions. They aid and abet the search for m-associated properties which turn ‘is the sun’ into a Juliet-friendly predicate. I have already argued that the way in which this is supposedly accomplished is flawed: Goodman's back-referential notion of exemplification simply cannot bear the weight that Stern puts on it in the context of metaphor. But what I am emphasizing here is that, for all it seems to be on the same wave-length as the semantic descent account, Stern's discussion of metaphors of exemplification does not depart radically from other familiar Content Sufficient accounts. Particularly striking evidence of this comes right at the end of Stern's book in a short section on ‘Non-linguistic Metaphors’. Having given a few examples of the construction ‘A is a metaphor for B’, where the expressions A and B pick out non-linguistic objects, he writes: ‘[I]f a necessary component of metaphorical interpretation is task-specific linguistic, or semantic knowledge, it should follow that only objects of language can themselves be metaphors. And if that is the case, all these non-linguistic symbols would be at best metaphorically, but not literally, metaphors’ (Stern 2000: 302).

Thus, where I have described the locution ‘is a metaphor for’ as the most direct possible evidence for the semantic descent account, Stern dismisses it as no better than a metaphorical way of speaking. (A strange put-down from an enthusiastic theorist of metaphor, but that is another issue.) To be sure, I have more than once insisted that qualification *on its own* is not metaphor, and that metaphor must begin in language. Stern and I share the intuition that the linguistic is a ‘necessary component’ in metaphor. But the semantic knowledge that Stern speaks of is fully visible in assertions fitting the schema ‘A is a metaphor for B’, since A and B are in fact linguistic expressions. However, while we agree about the starting point of metaphor being linguistic, our accounts branch sharply at the next stage. Where he is willing to allow objects as helpers in the quest for alternative meanings for the original expressions, I count them as predicates in their own right. And that is why the ‘A is a metaphor for B’ construction is a problem for him. He cannot see any place for alternative properties to be grafted on to the expressions in this construction, since both A and B seem simply to refer to non-linguistic objects. In contrast, I see the referentiality of A and B as just what we would expect. In ‘A is a metaphor for B’ we have an explicit appeal to the qualificational stage of metaphor, and this appeal requires that the object answering to A qualifies that answering to B.

As a final comment on the comparison between ‘Mthat’ and ‘↓ ... ↓’, let me briefly return to the model suggested by quotation marks. Stern's ‘Mthat’, like quotation, is intended to apply to any expression, but as we have seen this cannot be quite right, since the rule governing this operator speaks of ‘associated properties’, and this suggests that the expressions involved must be predicates. My ‘↓ ... ↓’ does not impose any such restriction on the expressions in its domain, but even so it is not

quite as liberal as quotation. Indeed, though I never said that ‘↓ ... ↓’ was exactly the same as quotation, it is important to be clear about the limitations of the analogy.

Quotation applies unrestrictedly to linguistic expressions, systematically altering their meanings by stripping away most, though not all, of their semantic functions, and replacing them with a special kind of referential function. Thus, while (7) *The cat is on the mat* contains both a phrase referring to a particular cat and another predicating a location for that cat, this: (8) ‘*The cat is on the mat*’, refers neither to a cat nor gives us predicative information about its location. Instead, we are to understand (8) as a device which in some sense refers to (7). I say ‘in some sense’ that it refers, and I suggested above that quotation does not strip away all of the semantic functions of (7), because as we have come to realize quotation is not quite a cut-and-dried referential device. It scarcely could be, given that adding the truth predicate to (8)—without removing the quotation marks—somehow reanimates the functions fulfilled by (7).¹⁷⁸

If the model of pure linguistic reference is that of a label or pointer, and nothing more, then quotation is not purely referential. But as philosophers have been discussing for the last hundred years, neither it seems are proper names, descriptions, demonstrative descriptions, or even, *pace* some remarks of Russell, ‘naked’ demonstratives. This is of course a can of worms, as is the proper story about quotation, but it seemed necessary to open the lid just a bit so as to show that my notion of semantic descent is in good company.

In applying the down-arrow notation to an expression—in marking the expression for semantic descent—we do not thereby create a device that is *purely* referential, nor do we get something that is exactly like quotation. Nonetheless, I think it perfectly reasonable to treat semantic descent as falling under the genus *reference*. In Chapter 4, I gave examples in which semantic descent conducted us from adjectival or predicative nouns to particular concrete objects, from verbal phrases to events and from whole sentences to states of affairs. It is being thus conducted that I regard as a form of reference, though I do not intend thereby to assimilate the relationship between these various phrases and the world to the species of reference that, for example, obtains between a proper name and its bearer.

Given that semantic descent is referential in the way just described, a difference between semantic descent and quotation shows up clearly. Whereas quotation has linguistic expressions in general in its domain, semantic descent only makes sense when applied to expressions that could in fact fulfil this more liberal referential

¹⁷⁸ There is also the complex phenomenon of mixed quotation, and, while lots that has been written about this would be congenial to the semantic descent analogy, especially given remarks in the next couple of paragraphs, I haven't had space to deal with it in the book.

function. Thus, whilst we can intelligibly put quotation marks around ‘or’ (as I just did), the construction ↓or↓ is of only doubtful sense.¹⁷⁹ Still, on one understanding of it, quotation exercises a referential function in taking us from the use of linguistic expressions to the mention of expressions so used. And semantic descent shares the same aspiration, whilst going in the opposite direction: beginning with a use of linguistic expressions, we move to the object or objects which we can count as the subject matter of those expressions.

5.2.5. Context

The combination of Stern's ‘Mthat’, which is intended as a general, metaphorical-meaning-inducing operator, with the importance he gives to ‘metaphors of exemplification’, suggests his account is structurally similar to mine. But these structural parallels do not stand up to close examination: Stern's ‘Mthat’ is not semantic descent, and his use of objects in metaphors of exemplification is radically different from mine. Moreover, as I have argued, there are irremediable defects in both these elements of Stern's account. It is as if Stern aimed in the direction of the semantic descent account, but didn't get close to it because of his adherence to the paradigm of meaning shift one finds in traditional Content Sufficient accounts. I do realize that this claim can seem both immodest and condescending, but it is certainly not intended to be either. Stern's valuable book presents an account which suffers from certain basic problems. However, if viewed from the perspective of my account and the arguments supporting it, one can see why these problems have arisen, even while recognizing what is right in Stern's account. Rather than counting as a denigration of Stern's account, I see this as a desirable kind of ecumenism, and as a further source of support for the semantic descent account.

This ecumenical theme is reinforced by the things that Stern says about the role of context. Context-dependency is a crucial feature of Stern's account—unsurprisingly given the title of his book—though there are fundamental elements of his discussion of context which are less than transparent. For example, while he mostly speaks of context as a set of presuppositions entertained by both speakers and hearers, at some places, he falls back on the more usual idea that context is a set of circumstances in which a linguistic act takes place. Yet these two conceptions do not fit together as easily as the shifts between them in *Metaphor in Context* would suggest.

Pursuing issues like this is well beyond my remit here. Instead, I should like to focus on a single underlying feature of the notion of context, one which deserves to be described as the *fulcrum* of context. What I have in mind is simply this: everyone agrees that contextual information is necessary, indeed crucial, for helping us (theorists) make sense of how we (speakers) make sense of linguistic acts, and nowhere is this more evident than in respect of metaphor. But for contextual information to be of any real use to theorists, there must be no doubt about the relation between

¹⁷⁹ That said, it has been suggested to me that one might well even treat the truth functions as objects in my sense. Think of: ‘Their relationship never had a chance: he was an and, she an or’.

given linguistic elements and contexts. It is not enough to say simply that there is some utterance and some context, and that we can better understand the former by appealing to the latter; instead, we must have some idea of how the context bears on specific features of the utterance. Moreover, this fulcrum or pivot of contextual information must be one that could be justified. Something must be said to explain *how* particular ingredients of context play the roles they do with respect to the relevant elements of the utterance.

A brief appeal to the index-set notion of context can illustrate this point. In normal circumstances, the assertion: (9) He delivered the parcel, creates no special problems of comprehension. As theorists, we say that this is largely because the normal circumstances—the context—includes information which makes it possible, even easy, for the hearer to know who is being spoken about, and when the action occurred. (For illustrative purposes, I am only focusing on pronoun reference and tense.) The context here could be described as a set of indices—subject and time(s)—which fill out the pronominal and temporal markers in (9). Note the fact that the relevant elements of (9) are described as pronominal and temporal *markers*; the pronouns ‘he’ and the suffix ‘-ed’ pretty much demand certain kinds of referent. Though the index-set can be arranged so as to yield up its elements in the order they are needed in the sentence, this is only to prevent confusion if there are a number of pronouns needing referents. The real basis for the link lies in our prior understanding of the relevant elements of (9), and each provides a fulcrum on which the weight of context is brought to bear.

Now the index-set idea of context certainly has its limitations, and I don't for a minute suggest that this rather formal, *semanticized* idea of context exhausts the importance of this notion. However, as just described, it illustrates my point about the role of context: any story about that role must make plain how contextual information is brought to bear on specific features of the utterance. It is not enough to say simply that elements in the context relate to elements in the sentence; we need to know something about the latter which actually leads us to the relevant features of context.

My worry about Stern's story—and the similar stories that have been told by many Content Sufficient theorists—is that this requirement tends to be overlooked. Recall in this regard the centrepiece of Stern's account, the rule governing the ‘Mthat’ operator: (Mthat) For every context c and for every expression ϕ , an occurrence of ‘Mthat [ϕ]’ in a sentence $S \dots$ in c (directly) expresses a set of properties P presupposed to be m -associated with ϕ in c such that the proposition $\langle \dots P \dots \rangle$ is either true or false in the circumstance c . We are told here that the context includes speaker-hearer presuppositions about properties m -associated with an expression in the relevant sentence. Previously expressed worries aside, this certainly seems to link context to elements of the

sentence: it is the context that supplies properties, and it these properties which hearers are seen as using in their dealings with \emptyset . Nonetheless, this appearance is deceptive. Because the 'Mthat' operator strips away the ordinary or literal meaning of \emptyset , nothing remains to be understood in the utterance which can point us to the item of the context we are meant to use. As an illustration, imagine an operator applied to 'he' in (9) which demanded that our ordinary way of understanding that pronoun was to be replaced, on a context-by-context basis, by some other way of understanding it. Since, so to speak, 'he' would no longer mean he, this would leave it mysterious how we would go about picking out that item in a given context which is linked to 'he'. My worry is that Stern's proposal leaves the link between context and utterance in no better shape.

It might be thought that, whatever we would think about the example of 'he', the following reasonable defence of Stern's notion could be mounted. As per the (Mthat)-rule, the expression \emptyset is understood to function predicatively in the original sentence. That is, it is understood as picking out properties which then take their place in the proposition expressed by the sentence in which \emptyset occurs. What the (Mthat)-rule does is to suppress the original fixed predicative meaning in favour of a variable one. That is, we are to imagine that, on a context-by-context basis, different properties are supplied by that very same predicate expression.

Does this work? Only superficially, and the weakness in this line of thought becomes clear in its final sentence. It would indeed be a defence of Stern's treatment of context if the properties it contains were supplied *by* the reconstrued predicate expression. But all we are in fact entitled to say is that the context supplies properties *to* that expression. Until the context does its thing, there is no reconstrual of \emptyset , so we cannot lean on such a reconstrual in getting from \emptyset to the appropriate item in the context (a property or properties). In fact, the only guidance we have in linking properties to 'Mthat- \emptyset ' is the clause in the (Mthat)-rule which insists that the result must be 'either true or false in the circumstance *c*'. But this is hopelessly weak. There are bound to be many properties which are m-associated with \emptyset , and make the resulting sentence either true or false in the relevant circumstance, without, however, producing an acceptable rendition of the metaphor.

A brief look at the role of context in my account should help clarify this criticism of Stern. When we apply the down-arrows of semantic descent to 'the sun' in Romeo's (R) we do not thereby strip this latter expression of ordinary meaning or convert it into an expression which is context-sensitive. Nor are we sent off on a search for various different properties, since the only relevant property turns out to be the obvious one of *being the sun*. Thus, the perfectly ordinary, fixed meaning of this expression, that is, the sun, is preserved throughout. However, when embedded in the predicate frame, 'is ...', the down-arrow expression yields a special kind of predicate, namely: 'is \downarrow the sun \downarrow '. This hybrid predicate—part word, part object—functions non-linguistically, or at least not purely linguistically. That is, it drafts in the sun to play a role that the sun, like any other object, is suited to play: the role of qualifying other objects. What constrains our understanding of the sun's role here? This is certainly where context is vital. As discussed at the end of the previous

chapter, there are three main types of constraint which can be thought of as contextual. First, there is the particular way of referring to the sun that Romeo used, and the immediate linguistic context of that reference. Secondly, there is the broader context that includes those things we know and think about the sun—its cultural significance, as I have described it. And finally there is the kind of general context within which the linguistic act takes place. The first two kinds of context are specific to the qualificational use of the object got by semantic descent, but this third is not in essence different from the kinds of general context appealed to in helping us understand linguistic acts not specifically metaphorical. Very roughly, general context includes grasp of the setting within which the act is effected, where this includes knowledge about the subject of the act and its author.

In my account, appeals to these kinds of context, especially the first two, are not for the purpose of finding some property or properties linked to the expression ‘is the sun’ which might then plausibly apply to Juliet. This contextual material is brought to bear, not on any expression, but on an object, and its point is that of helping us with the sun's qualification of Juliet. Here we have a perfectly determinate object (the sun), brought to our attention by the unsurprisingly suitable expression (‘the sun’), and it is this object which serves as the fulcrum or pivot of contextual information. Just as ‘he’ in (9) demands certain kinds of referent from the surrounding context, so it is the sun (not ‘the sun’) which brings down on itself the weight of specific features of context. The more we know about the sun in the relevant context, the more likely we are to understand Romeo's (R).

As noted earlier, Stern's story about metaphors of exemplification can sound very much like semantic descent. But it is not the same story. For Stern, the purpose of context in respect of metaphors of exemplification is that of finding some suitable property or properties which we can then link to ‘the sun’. Admittedly, in respect of metaphors of exemplification, this search does actually call on the services of the sun. But this good feature is diluted by the fact that what is asked of the sun is the production of properties that it exemplifies.

One final note: I can imagine someone impatient with the general line of argument just presented. The complaint would run: you go on about the sun as a qualifier of Juliet, Stern looks for a property that the sun exemplifies which also applies to Juliet. In the end, is there really much difference between these two proposals?

The idea that at bottom metaphor involves a search for special, or in some way non-standard, properties has surfaced several times before in this book. After considering White's account of metaphor, one which also leans towards property-involvement, I shall return to it. But the short answer to the above question is ‘yes, there is an enormous difference’. My account is most certainly not merely a notational variant of these property-based accounts.

5.3. The Conflated Sentences Account

Roger White's *The Structure of Metaphor* (1996) is one of the most valuable works on the subject in recent years. One of its contributions is the robust unmasking

of defects in many other accounts of metaphor. Black, Davidson, Searle, and many other theorists are taken severely to task for making claims that are simply false, or are implausibly promissory. Yet even more valuable are the many sensitive observations that White makes about the phenomena of metaphor (as he calls them), phenomena such as bifurcation, polarity, and mixed as well as extended metaphor. However, having acknowledged many of these critical and observational details in earlier chapters, here I propose only to consider White's own account of metaphor. This will be a considerably less complicated business than my discussion of Stern because White's account is not built on a detailed logical substrate. In fact, the outline of the account is rather straightforward. White sums up the core of his view in the following admirably succinct way: 'A metaphor is a sentence that may be regarded as a sentence that has arisen from the conflation of two other, grammatically analogous, sentences' (White 1996: 106). The idea that this is only the *core* of White's account must be taken seriously; the full account involves an important reconstrual of this initial claim. Still, the reconstrual depends on this initial statement, so let me first illustrate White's idea of metaphor as a 'conflation' of sentences, using one of his own examples.

5.3.1. A case of two sentences

In Act 4 of *Othello*, at Cassio's appearance, Iago says to himself: Here he comes. As he shall smile, Othello shall go mad; And his unbookish jealousy must [construe] Poor Cassio's smiles, gestures and light behaviours Quite in the wrong.¹⁸⁰ The sentence-sized element that White extracts from this for discussion is: (10) His unbookish jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong. That the unit which one identifies as a metaphor is sentential, rather than word or phrase size, is something White insists on. Indeed, he pours scorn throughout the book on accounts which begin by assuming that metaphors are particular ways that words or phrases are used. White is of course perfectly right to note the violence one would do to this and similar metaphors by picking and choosing this or that metaphorical word; obviously enough, there is a metaphor theme that runs through the whole of (10). I will revisit the question of the size of metaphors later, but note here only that this question is not an open one for White. As the above citation makes plain, *all* metaphors are understood as a kind of mixing or conflation of

¹⁸⁰ *The Oxford Shakespeare* (Shakespeare 1988: 842). This text has 'conster' instead of 'construe', but then again White uses a 19th-century text that has, among other oddities 'Ielousie' for 'jealousy', so I have compromised.

sentences. In the particular case of (10), two sentences which are candidates for such conflation are: (10a) His uncultured jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong. (10b) The unbookish schoolboy must construe the *Iliad* quite in the wrong. But it is important that they are not the only pair relevant to (10). White's proposal for coming up with candidate sentences begins by finding in (10) the following pair of open sentences: (10c) His *X* jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong. (10d) Unbookish *X* must construe *Y* quite in the wrong. These open sentences result from replacing certain expressions in (10) with variables: (10c) is the result of replacing what White thinks of as an item in the primary vocabulary in (10), and (10d) results from replacing two items of the secondary vocabulary.

Note that this distinction between vocabularies is not to be treated as tantamount to the metaphorical/literal distinction; White insists that both vocabularies are literal, though each is appropriate to a different subject matter. Thus, in moving from the open sentence (10c) to the closed (10a), we must replace any variable by a term which belongs with the rest of the vocabulary used in (10c). In particular, since (10c) is essentially a description of how a human being's strong emotion can influence the interpretation of another's behaviour, the replacement for *X* must fit that scenario. 'Uncultured' is one possible filling; perhaps 'unsocialized' or 'crude' would be others. Similarly, (10d) calls on vocabulary appropriate to the untrained efforts of someone trying to interpret something like a text. Here White suggests replacing *X* with 'schoolboy' and *Y* with 'the *Iliad*', though endless variations on this same theme are possible.

Once we have moved from the two open sentences to the pair of closed sentences (10a) and (10b), we work backwards, as it were, to imagine how Shakespeare came up with (10) in the first place. Here is White on how this is done: In constructing such a metaphor out of two such sentences ... the constructor makes a series of choices between the pairs of words or phrases offered by the primary [10a] and secondary [10b] sentences. The only restrictions upon the choices made is the consideration that, if the hearers of the metaphor are to understand the metaphor correctly, they must be able to reconstruct appropriate primary and secondary sentences for the metaphor. (White 1996: 79) Each of (10a) and (10b) are, as White insists, literal descriptions of different situations. What Shakespeare did—or what we could imagine he did—was to make a judicious set of choices in deleting various items from each of these sentences, and putting the results together—conflating them—so as to create the well-formed sentence (10). It is suggested above that the *only* restriction on choice and conflation is the possibility of hearers being able to recover the primary and secondary sentences. But of course there must have been lots of other factors that led Shakespeare to the

particular metaphor (10), and White's book contains many detailed discussions which touch on these other factors, in this and in many other cases. So in fairness to White, we should read the above as offering a minimum or necessary condition of intelligibility, and not the whole truth about metaphor construction.

Accepting what White says about (10) is not difficult, but of course the story about the conflation of two sentences is not yet an account of metaphor. For that we need, minimally, to be told some story about the point of this kind of conflation. White fills in that story this way: One of these sentences, the primary sentence for the metaphor, is a sentence that would give a literal description of the actual situation. The other sentence would give a description of a situation with which the metaphor invites us to compare the actual situation. As a result of such conflation, we are invited to explore a network of similarities and dissimilarities between the two situations, and see the one situation in terms of the other situation, to see it as if it were the other situation. (White 1996: 80) As noted above, the story about the conflation of two sentences is only the core of White's account, its real but only initial starting point. Not wanting to be stuck with an apparatus that can make it seem required, case-by-case, to find two fully determinate sentences for each metaphor, White comes to treat the talk of sentences as eliminable. It is not that (10) really does result from sentences like (10a) and (10b); after all, what Shakespeare makes Iago say is just one sentence, one metaphor, and it is this sentence that an account must explain. The final twist to White's account is well summed up by the title of the section in which he explains it: 'The Metaphorical Sentence as a "Duck-Rabbit"' (White 1996: 113). In more detail: [The metaphorical sentence] is a sentence that may simultaneously be regarded as presenting two different situations; looked at one way, it describes the actual situation, and looked at the other way, an hypothetical situation with which that situation is being compared. These two ways of looking at the sentence arise by switching some of the words it contains from their normal, 'proper', use to a use as dummy names: switch the words of the secondary vocabulary from their proper use to a use as dummy names, and the sentence becomes a description of the actual situation ... switch the words of the primary vocabulary, and it becomes a description of a different situation ... (White 1996: 115) Exactly how pulling this duck-rabbit out of a hat reduces any implausibility in the two-sentence treatment is not clear to me. But there is no need to pursue this here, as my sole interest is in the account of metaphor that White offers *after* he has told his two-sentence/duck-rabbit story. As both of the passages above make clear, the account of metaphor White offers is a variant of those based on comparison and similarity. I say, 'a variant' because there are features of the two-sentence/duck-rabbit story which make White's appeal to similarity somewhat different from others. Before explaining how it differs, I need to revisit some of my criticisms of appeals to similarity in accounts of metaphor.

5.3.2. Similarity and property-sharing

Even a minimal familiarity with the literature would reveal how often similarity ends up figuring somewhere in accounts of

metaphor. In some, the similarity idea is centre stage. Metaphors are treated as elliptical similes, and these latter are understood as judgements of similarity; by a simple inference, we end up with an account of metaphor as similarity. Other accounts feature similarity but somewhat less centrally. As we have seen, underlying Stern's demonstrative account is a story about the context-sensitive selection of properties associated with the original metaphor expression. As a result of this selection, the metaphor expression acquires a different meaning (i.e. character). Similarity plays a crucial role in this, even though Stern is careful to insist that similarity is not what the metaphor sentence asserts. He says that similarity is located 'in the context of the metaphor rather than as part of its content' (Stern 2000: 147). So, Romeo is not *saying* that Juliet is similar to the sun, though it is a feature of our understanding of his claim (and his understanding of it) that it is *true* that Juliet and the sun share a property or properties, and also true therefore that they are similar. And, Stern's account aside, some notion of similarity figures in Content Sufficient accounts like those of Black, Beardsley, Kittay, and, despite his strictures on appeals to similarity, Goodman. So White's appeal to comparison and similarity has a long pedigree.

Despite the way I have associated them, the relationship between similarity and the idea of property-sharing is not straightforward. While it seems a truism that if A and B share a property, they are to that extent similar, not everyone thinks it obvious that the converse holds. For example, Fogelin (1988) insists that two things can be similar even though they don't share any relevant property. This idea, which he describes as 'brute' similarity, is offered as part of a defence against a certain kind of attack on his view. Because the fortunes of brute similarity are important to what follows, I need to say more about this attack and defence.

Central to the attack is an observation that I have myself made at various points: in any typical case of metaphor, it is implausible that the metaphor subject and the metaphor vehicle actually share any property, except in either a trivial or equivocal sense. We may be tempted to say that, by asserting (R), Romeo claims each of Juliet and the sun to be *warm* and *necessary for life*. But if we are honest we will see that these properties are quite different in respect of the sun and Juliet. Thus, we only equivocate if we claim the members of each pair to be the same, and we rely for any similarity we are tempted to find on the very thing we had set out to explain, namely metaphor. In the particular case of Fogelin's simile view, the attack concludes: unless there is a property shared by metaphor subject and vehicle, there is no independent justification for counting them as similar. Given this apparently anti-similarity conclusion, it seems impossible to maintain the simile view of metaphor.¹⁸¹ However, if Fogelin is right in thinking that there can be similarity without property-sharing, then not only does this attack on the simile view fail, so might this argument against similarity accounts of metaphor generally. Everything therefore depends on whether brute similarity is a viable notion.

Fogelin illustrates brute similarity with two examples, the first of which is that two different shades of blue can be judged similar, even without there being a

¹⁸¹ Fogelin is responding to an argument he finds in Searle 1979/1993.

property that they share. This example is not convincing, since it is easy to imagine lots of properties of two colour samples that explain their similarity: a good judge of colour might find two shades of blue similar because they each 'contain red' or are each 'pastel' etc. Moreover, the colour example is not really pertinent to metaphor, since shades of blue are items in the same category, whereas in metaphor any similarity will cross categories. However, Fogelin offers a second example, one which supposedly illustrates brute, cross-categorical similarity: 'Mathematical proofs, recipes, designs, the way to a friend's house, the instructions for assembling a lawn chair, can all be simple' (Fogelin 1988: 47–8). However, the problem with this example seems obvious. I could imagine someone objecting that, in describing this case, Fogelin actually gives us the property that each of these different things share, namely *simplicity*. So, how can he regard this as a good case of brute similarity?

Fogelin doesn't follow up on his example, but there is certainly something he could have said to head off this objection. Indeed, he could have appealed to something like the observation that served as the premise in my argument against the appeal to similarity in metaphor. Simplicity is not the *same property* in respect of mathematical proofs, recipes, designs, etc. But because we find resemblances amongst certain properties of these things—one wants to say amongst these different simplicities—we find it natural to say that they are all instances of simplicity. In effect, our attribution of simplicity to all of the instances follows from, and does not ground, our antecedent judgement of brute similarity.

That this defence of similarity uses the same premise I appealed to in attacking similarity underlines the point I made earlier: my argument seems only to be effective if one begins by assuming that similarity requires property-sharing. In so far as we accept Fogelin's contention that similarity can be brute, that argument is blocked. Does this leave the way open for appeals to similarity in the explanation of metaphor? That is a further question, one which can only be answered by looking more closely at the very idea of brute similarity.

Fogelin's two examples of brute similarity are not robust enough to bear the weight he puts on them; on their own they would not be enough to vindicate either the simile view or any more general appeal to similarity in the explanation of metaphor. The shades of colour example is not convincing at all and, while I have defended the cross-categorical example against the obvious objection, it is easy to imagine stronger objections. Still, I think one can conjure up more convincing cases of supposedly brute similarity.

Both of Fogelin's examples involve the similarity of relatively concrete, nameable objects: one shade is said to be brutally similar to a second, and a recipe is said to be brutally similar to a design. This kind of example invites, and certainly cannot close off, a certain obvious kind of objection. Fogelin insists that the similarity is in each case brute, but the objector might, with some justification, think that this is merely defeatism, that if we think hard enough, we will be able to find properties of each thing which ground the similarity judgement. What would serve Fogelin's purposes better would be a kind of similarity judgement in which one sees at the outset the

fruitlessness of the search for further properties. My suggestion is that we are more likely to find such cases by looking at the things we say when comparing situations, events, and actions—the very kinds of comparison that figure so centrally in White's account of metaphor.

Before getting to White's examples, first consider one of mine from Chapter 4: Swerving at the last moment to avoid innocent bystanders, his argument came to halt. Underlying its surface syntactic complexity, this sentence conjures up two situations: one a dangerous swerving and stopping of some vehicle under minimal control (perhaps the action of a madcap bicyclist), and the other the presentation of a rather abrupt argument (perhaps a philosophical one).¹⁸² By my lights, the recovery of these situations is due to the process of semantic descent, and metaphor works by having one—the madcap swerving—as a predicate of the other—the lame argument. However, for present purposes, forget about my account of metaphor, and think only of the idea of comparing these two situations. Are they similar? Do they share any property that could explain any similarity we find? Leaving aside my own inclination to answer both of these negatively, I can well imagine someone answering the first, and only the first, affirmatively. Here is how this might go: the bicycle event is like the argument event, but this likeness is somehow based on our overall assessment of, or reaction to, each of them; it is certainly not grounded on the presence of properties of one actually shared by the other. This is not to deny that the situations might be said to have properties in common: both involve someone doing something without exercising control, and both involve an abrupt ending. But, even aside from the familiar stretching of words that these supposed properties require—perhaps stretching to the point of equivocation—the plain fact is that they do not themselves explain the similarity. There are endless other kinds of event that involve lack of control and sudden endings, but which would not be judged similar to either or both of the events at issue here.

The suspicion is that we are tempted to find shared properties as a result of the similarity, so this cannot be an explanation of it. Speaking with Fogelin here, one would say that there is something brutally similar about the two, and it is this brute similarity which encourages us to invest more in any alleged common property than is merited. Speaking for myself, I would add that it is our finding aptness in the metaphor which tempts us into saying that the situations are brutally similar; the brute similarity judgements, if made at all, follow from, and do not explain, the metaphor. But I shall leave this subversive thought on one side until after considering White's example.

White is perfectly happy to find comparison, that is judgements of similarity and dissimilarity, at the root of metaphor. But the issue at present is the *basis* of these

¹⁸² The sentence in isolation isn't about bicycles or philosophical arguments, but one can easily imagine a larger context which would incline one to these determinacies.

similarity judgements, so let's take a fresh look at the *Othello* example. According to White, in uttering: (10) His unbookish jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong, Shakespeare's Iago conflates two sentences, or, more in keeping with White's final view, he uses one sentence that is intelligible as a description, either of this situation: His unsocialized jealousy will lead him to misjudge poor Cassio's smiles, gestures and light behaviours, or this one: The unbookish schoolboy is bound to construe the *Iliad* quite in the wrong.¹⁸³ White's account requires him to find these two situations importantly similar, but would he also accept their sharing some property or properties that explains the similarity? I very much doubt it, and for good reason. It is extremely unlikely that a situation in which some individual's state of mind influences his view of others' behaviour has substantial properties in common with an untrained schoolboy's attempts at understanding the *Iliad*. No doubt, we could manufacture some property they have in common, for example, we could say they both involve interpretation and mistake. But, as already noted in connection with my example, these scarcely *explain* the similarity. First, many situations very dissimilar to those above involve interpretation and mistake. Second, the interpretation and mistake specific to each situation are quite different properties. As in the previous example, what makes us even think they might be the same is the similarity of the situations. So White's example, like mine, is one in which we can compare situations, and make judgements about their similarities and dissimilarities, without in the least thinking that these judgements are explained by, or grounded in, independently accessible properties that the situations have in common.

For this reason, these examples lend support to Fogelin's notion of brute similarity, though they do so at a price. Let me explain.

Fogelin appeals to similarity to explain metaphor. This seems to be undermined by the fact that, in metaphor, the relevant items do not share properties, and are for that reason not similar. Fogelin counters by insisting that there is a kind of brute similarity which can obtain in the absence of property-sharing, and that it is appeal to this notion that underpins his account of metaphor. However, since it seems that the really plausible cases of brute similarity are those which involve metaphor explicitly or implicitly, there is something suspicious about Fogelin's appeal to this kind of similarity in his explanation of metaphor. Unless something further can be said about brute similarity—something which is itself independent of metaphor—we have no reason to accept the final stage of Fogelin's defence.

¹⁸³ I described the two situations slightly differently from White. As will be clear, these new versions (quite licitly) serve my purposes better, and I think they also better serve White's purposes.

But there is, or at least seems to be, something we can say, something White alludes to in the definitive statement of his account. Of his idea of metaphor as the conflation of two sentences, each describing, non-figuratively, a possible situation, White writes (as already noted): ‘As a result of such conflation, we are invited to explore a network of similarities and dissimilarities between the two situations, and see the one situation in terms of the other situation . . . , to see it as if it were the other situation.’ White here connects similarity with metaphor—the passage tells us what we are invited to do by metaphor—but he does not thereby suggest that we can explain metaphor merely as similarity. For he goes on to suggest that the presumably special kind of similarity we have in metaphor is itself explained by the further idea of seeing one thing as another. If this latter notion does help us to understand the special kind of similarity we have in metaphor, then this would address the circularity that threatens Fogelin's account.¹⁸⁴ But, from Fogelin's point of view, the cost of this is high. For, while we might in this way vindicate the idea of a special kind of similarity in metaphor, we would have left behind the simile thesis. What would be pivotal in metaphor would not be similarity, but rather the notion of seeing one thing as another.

White's appeal to the notion of seeing-as creates a big problem for me too. If, as White claims, it could form the basis of a satisfactory explanation of metaphor (as well as of similarity), it would weaken the case for there being a special kind of predication in metaphor. However, though it will take some showing, so far from this, I think that any reasonable attempt to understand the idea of seeing-as leads us straight back to a key element in my view.

5.3.3. Seeing one thing as another

Appealing to the notion of seeing-as in respect of metaphor is not original to White; Davidson says positive things about it, as do Stern and many other Content Sufficient theorists. Yet, for all that it figures in discussions of metaphor, it tends to be under-described. There seems to be an unstated assumption that we know enough about this notion from the ordinary contexts in which it occurs to understand its significance for metaphor.¹⁸⁵ However, I don't think this is satisfactory.

What do we mean when we say that one thing is seen as another? The notion of seeing-as figures centrally in disputes about perception, and it has generated a massive literature. Still, the point that I need to make is off to the side of issues about perception, so I can make it without treading on the minefields in that literature.

Consider two simple cases. First, a subject is shown the Müller-Lyer figure, and claims:(11) I see the lines as unequal.

¹⁸⁴ White doesn't speak of brute similarity, but, given the things said above about the *Othello* example, I suspect that he would accept that the similarity between the relevant situations is brute in this sense: it cannot be explained by any independently accessible shared properties. In this regard, look at what he says about common properties on p. 243.

¹⁸⁵ White does say quite a bit about seeing-as and metaphor, but he doesn't seem to me to ask the right questions about it.

Second, you see an object in the distance, and claim:(12) I see it as a pony. It is tempting to think that the claims made in each case could as well have been, respectively:(11') I see the lines are unequal,(12') I see the distant object is a pony. It is tempting, that is, to treat the seeing-as idiom as including straight predication, and, as will be argued, I think it is right to be so tempted, since I shall suggest that it explains a lot about the appeal of the 'seeing-as' idiom in respect of metaphor. But there is an obvious obstacle to this treatment of 'seeing-as'.

Each of (11') and (12') implies the embedded claims, respectively, that the lines are unequal, and that the object is a pony, but neither (11) nor (12) have any such implication. Moreover, each of (11') and (12') implies that the speaker *believes* these embedded sentences, even though neither (11) nor (12) have that implication. Someone may well think that the lines are equal, and still insist on (11), or may well know the object to be a tree-stump, and insist on (12).

Both of these difficulties come from a tendency to think of 'see (that) ...' as factive. Because of this we hear (11') and (12') as implying the truth of the embedded predications, and from this, together with the standard idea that a speaker believes the obvious implications of what he asserts, we infer that one speaker believes that the lines are unequal, and the other that the object is a pony. Both of these unfortunate implications can be dealt with by making sure that the factive character of 'sees' is suppressed. One way we might try to do this would be to rewrite the offending sentences as:(13) It appears to me that the lines are unequal,(14) It appears to me that the distant object is a pony. Still, this is not a perfect solution, since, depending on how one takes the 'it appears to me' idiom, each of these can suggest either that their respective speakers believe the embedded sentences are true, or believe they are false. Perhaps better would be:(15) Appearances suggest to me that the lines are unequal,(16) Appearances suggest to me that the object is a pony, but there is no real need to be too fussy about finding a precisely correct form of words, since, as noted above, my ultimate interest in the 'seeing-as' idiom is outside the perceptual context. For that use, we could as well take ourselves to be dealing with:(17) I conceive the lines as unequal,(18) I conceive the distant object as a pony,

and these really do seem equivalent to the clearly predicative:(19) I conceive the lines to be unequal,(20) I conceive the distant object to be a pony.The upshot is this: the idea of seeing A as B, and certainly the idea of thinking of, or conceiving, A as B, can be understood as including the predication A is B, so long, that is, as we are careful to remove the assertoric suggestion that comes with the latter. This is because the real point of 'as' in 'seeing A as B' is to force us to step back from the outright assertion we have in 'seeing A is B'. However, so long as we can achieve this step back, say by using the non-factive 'appearances suggest', or 'conceive', we can reinstate the predication that is at the heart of the seeing-as idiom.

The above discussion concerns straightforwardly literal uses of 'seeing (or conceiving) A as B'. But what about the use of this phrase in connection with metaphor? Surely, the idea that Romeo's (R) is an invitation to see Juliet *as* the sun cannot be treated as an invitation non-factively to see or conceive that Juliet *is* the sun. To be told that Romeo is merely conceiving Juliet is the sun, not asserting it, is of no help given that we don't understand this predication in the first place. Matters are even worse in respect of White's example. He claims that the metaphorical:(10) His unbookish jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong,is a conflation of these two sentences:(10a) His uncultured jealousy must construe poor Cassio's smiles, gestures and light behaviours quite in the wrong,(10b) The unbookish schoolboy must construe the *Iliad* quite in the wrong,and that we should understand (10) as requiring us to see the situation described by (10a) *as* the situation described by (10b). But, if I am right about seeing-as, then this would also require us to conceive that the situation described by (10a) *is* the situation described by (10b). But it is surely, if anything, more difficult to conceive this than it is to conceive that Juliet is the sun.

Faced with these difficulties, one option would be to deny that seeing-as in connection with metaphor works like seeing-as in literal cases. In the latter, 'as' includes 'is', but includes also some way of dampening down the assertoric implications of 'is'. Perhaps the kind of comparison one finds in metaphor calls upon a notion of seeing-as that is independent of this predicative treatment; perhaps this notion of seeing-as is *sui generis*. (That would certainly explain the reticence of many writers to say more about it.)

However, this option seems desperate, as well as unhelpful. As noted, seeing-as must be treated with great care, but, given that it seems to function in much the same way across contexts which range from the narrowly perceptual to those in which conception rather than perception is at issue, we need to be given some good reason for its sudden change of character in the context of metaphor. It smacks of desperation to find a sudden change in this idiom merely because the usual way of

understanding it creates problems when applied to metaphor. Also, insisting on a *sui generis* construal of seeing-as fuels the suspicion that we can learn nothing about metaphor by appealing to the idea of seeing one thing as another. The point of any such appeal is to cast light on metaphor, but a *sui generis* idea of seeing Juliet as the sun is no less mysterious than the original metaphor. We could have got to the same place by insisting that the metaphor predication, Juliet is the sun, is itself *sui generis*, and not like any ordinary literal predication.

There is however something much more satisfying that we can say about all this. My few remarks about the seeing-as idiom suggest that it includes predication, albeit with the epistemic consequences of this highly qualified. This creates problems for those who seek help with metaphor from this idiom: we are either led straight back to the very predication that creates the problem of metaphor in the first place, or we go down what seems the blind alley of treating metaphor predication as itself *sui generis*. Predication seems the key to all of these difficulties and, precisely because of what it says about predication, the semantic descent account has a key role to play here.

According to this account, we do indeed find predication in the claim that Juliet is *seen as* the sun, the same predication as we have in the claim that Juliet *is* the sun. However, what makes all this possible, as well as unmythical, is the fact that the vehicle of this predication is not the words 'is the sun', but the sun itself. Recognizing that objects as well as words can function predicatively (i.e. qualificational), we can preserve the univocity of the seeing-as idiom, while at the same time explaining its special importance in respect of metaphor. Seeing-as remains univocal because it does indeed include predication, even though, in the case of metaphor, it is not linguistic predication. Following on from this, we can also appreciate why so many have been tempted to appeal to the seeing-as idiom when confronted by metaphor. Let me spell out this second point.

Disappointed that the words 'is the sun' fail to work straightforwardly as a predicate of Juliet, one can see the attraction of another, apparently less problematic, way of using virtually these same words. Instead of having to conceive that Juliet is the sun, we need only conceive of Juliet as the sun. As noted above, this turns out to be unsatisfactory, precisely because it doesn't take us far enough away from the predication that had disappointed us in the first place. But, given a determination to stick as closely as possible to the senses of the words in the original predicate expression, this move can seem mandatory.

White insists that the subject-predicate metaphor form has distorted virtually all accounts of metaphor, and he thinks that more realistically complicated metaphors should be the focus of our ruminations. But the above point applies just as well to his conflated sentence account. What in the case of the Romeo example is the determination to preserve the senses of the words in the predicate 'is the sun', in his account comes out as a determination to preserve the dual vocabularies whose conflation is the origin of the metaphor. And this is made possible by his insisting that we see the situation characterized in the one vocabulary as the situation characterized in the other. Still, while it is situations that are compared and contrasted,

the focus of White's explanation of metaphor is resolutely linguistic. In connection with Sonnet 65 (Shakespeare 1988: 759) O how shall summer's honey breath hold out, Against the wrackful siege of battering days, When rocks impregnable are not so stout, Nor gates of steel so strong but time decays she writes: What we have is a sentence which, in the way I have been outlining, can be read *both* as a description of the destructive effects of Time, *and* as a description of the use of a battering ram. Shakespeare has *superimposed* a description of the action of a battering ram upon a description of temporal processes. In so doing, he has used words which properly are used in the description of the battering ram, as names for the action of time, and, *in this way*, is talking of time as if it were wielding a battering ram. (White 1996: 117, his italics) This passage invites us to find metaphorical effect in the interaction of *descriptions*; White's conclusion even transposes the seeing-as idiom into 'talking-as'. In thinking of the metaphor this way, we can keep the words of the original, even though their original predicative brief cannot be fulfilled.¹⁸⁶

This picture of metaphor is in stark contrast to the idea that what Shakespeare is doing in the above is using one situation as a non-linguistic predicate of another. Or, that in the other case, Romeo is using the sun as a non-linguistic predicate of Juliet. I have argued that seeing-as, or talking-as, does not really take us far enough from the predication that is either on the surface, or just below, any metaphor.

But there is a second problem with it. Though I have tried my best to explain its attractions, many accounts based on comparison/brute similarity/seeing-as/etc. make metaphor implausibly weak. Romeo does seem to be saying something about Juliet, Iago does make an assertion about Othello. Yet these assertions are somehow weakened into invitations to compare and contrast. In contrast, with semantic descent we can have our cake and eat it. Romeo is indeed making a subject-predicate assertion about Juliet, albeit one that calls on the kind of non-linguistic predication I call 'qualification'. Nonetheless, given the way the seeing-as idiom works, there is nothing to stop our also thinking of Romeo's utterance as an invitation to see Juliet as the sun. Remember that though we do not want an assertion of 'seeing A as B' to commit us to the truth of 'A is B', or to the latter's being believed by the speaker, neither of these things is actually ruled out by the seeing-as idiom. That is, on a case-by-case basis, it could be true that someone who says 'I see A as B' is in fact also asserting that 'A is B'. This is of course not something that White, Fogelin, or any similarity theorist can allow, simply because their accounts cannot make sense of the predicational claim. But the semantic descent account, in giving us a way to understand the predication, not only preserves the relationship between seeing-as

¹⁸⁶ The sentence from Sonnet 65 is not of subject-predicate form, so what I say about the predicative brief of the original words is not wholly accurate. I shall deal (briefly) with this and related issues below in subsection 5.3.4.

and seeing-is, it also allows, when appropriate, for epistemically stronger claims than merely seeing-as. And in most cases of metaphor, this stronger form is appropriate. Thus, while Romeo might well be inviting us to compare Juliet to the sun—to see Juliet as the sun—he is also asserting that she is the sun. In other cases, one of which is about to be discussed, assertion might not be appropriate, and comparison would suffice. But, while the semantic descent account has resources to deal with both, accounts like White's do not, and are to that extent unacceptable.

5.3.4. The conflation of sentences?

White's account of metaphor depends, at bottom, on similarity. To be sure, the appeal to similarity is not as direct as that in Fogelin's simile account, and White does try to *explain* similarity by calling on the idea of seeing of one thing as another. But these do not save the view because: (i) as I have argued in several places, any plausible notion of similarity in respect of metaphor depends upon metaphor rather than the other way around; (ii) the predicational aspirations of metaphors cannot be avoided by appealing to the idea of seeing-as; and (iii) we should respect these aspirations, since we otherwise cannot explain the assertoric import of many metaphors.

Still, there is a niggling point that might have struck you in respect of my treatment of White's view, and, even if it didn't yet strike you, honesty requires me to spell it out. In the examples from *Othello* and from Sonnet 65, while there seems to be an underlying comparison of situations, the speaker in each case does not seem to be asserting that the one situation *is* the other. Whereas Romeo does both see Juliet as the sun, and say that she is the sun, Iago might well invite us to see the effects of jealousy as like the efforts of a hopeless pupil, but he never says that the jealousy situation is the pupil one. So, where does that leave my carefully constructed trap for White? I went through a lot of trouble to conclude in section 5.3.3 that it was an advantage of my account that it could make sense both of seeing-as and seeing-is. However, a closer look seems to show that, as far as White's examples are concerned, we might well not want to make sense of the seeing-is form.

You might be surprised that I describe this as a 'niggle', given that it seems to undermine my criticism of White. However, once things are spelt out a little more clearly, you should see that, when care is taken over certain syntactical matters, the difficulty disappears. Following the herd, I think of philosophical difficulties which turn on syntax as niggles, but, pursuing this one is worth the effort because it will lead to an interesting take on the relationship between White's account of metaphor and mine.

White adamantly insists that metaphors are *sentences* that are conflated, and many of his examples are of quite long sentences in which metaphor effects are, to use his term, extended. For example, in the Sonnet 65 example: O how shall summer's honey breath hold out, Against the wrackful siege of battering days, When rocks impregnable are not so stout, Nor gates of steel so strong but time decays,

there is a thread of metaphor running from ‘hold out’ to ‘siege’ to ‘battering’ and it is this thread, rather than any one element in it, which White regards as crucial to the comparison with the unmetaphorical and ordinary decaying effects of time. With all this, I certainly agree, but I do not think we can conclude from this that comparison is all there is to it—that predication is nowhere to be seen. It will help here to recall what was said about an example in Chapter 4:(21) When questioned, he offered his usual soap-bubble reason for what he had done. Though this assertion is not at bottom a metaphor, it contains one—the sentence constituent, ‘soap-bubble reason’—which is crucial to understanding the assertion. Now, a comparativist would claim that the metaphor here involves the comparison of a reason for action with a soap-bubble, that we are here invited to see the reason as a soap-bubble. However, aside from any difficulties we might have in explaining such a wildly cross-categorical comparison, there is something too weak about all this. Surely, one feels, there is more to the juxtaposition of these two things than an invitation to compare. This feeling, along with the urge to see comparison, can be easily explained by the semantic descent treatment of (21).

Adjective-noun sentence constituents like ‘stone house’ are taken in our stride—they are ways of saying of the house that it is made of stone—and there is nothing to stop us treating ‘soap-bubble reason’ in the same way. Once we recognize that an object—in this instance, a soap-bubble—can function as a non-linguistic predicate, it is possible to see the constituent in (21) not as a comparison, but as a predication. This doesn't mean that there isn't also some kind of comparison lurking here. For, when the constituent is understood as a predication, we can also see why we might further think in terms of both comparison and seeing-as. Given that the soap-bubble is a qualifier of this particular reason, the idea of comparing the two becomes intelligible, as does the idea of seeing this particular reason as a soap-bubble. (Think here of the fact that, given the truth of the predication, the class of things which are reasons will actually include the particular one that is qualified by the non-linguistically ‘adjectival’ soap-bubble.)

Against this background, recall this related sentence:(22) When questioned, he offered his usual soap-bubble reason for what he had done, a reason which burst as soon as the detective pricked it by citing several witnesses' statements. Here there is a string—a veritable thread—of metaphors. White insists that cases like this cannot be handled as if they involved word-based metaphors, occurring one at a time, and surely he is right. But there is a way of keeping hold of the thread that doesn't involve his apparatus of sentence conflation. Semantic descent takes us from ‘soap-bubble’ to a soap-bubble, and the latter is employed predicatively. However, once called upon, this object can also serve to encourage semantic descent elsewhere in the sentence, and to link the resulting predicative uses.

What lies behind my proposal for metaphors that thread through this sentence is knowledge about soap-bubbles: we know what they are like, we know that they

burst when pricked. Put into a phrase as I just did, it can look as though this knowledge vindicates White's story about sentence conflation. But this would be hasty: do we really want to consider (22) a conflation of these two sentences: (22a) When questioned, he offered his usual poor reason for what he had done, a reason which lapsed as soon as the detective queried it by citing several witnesses' statements, and: (22b) Soap-bubbles burst when pricked? Given the alternative story made available by the semantic descent account, the lack of balance between these two candidates makes the whole idea of sentence conflation far-fetched. (I simply cannot think of a way of beefing up (22b) in the fashion required to see some balance with (22a).) Indeed, we wouldn't have the first idea of how to bring these sentences into line in readiness for conflation, in the absence of the link between 'soap-bubble' and 'poor'. But any story about why knowing this is so crucial, is most of the way towards the kind of linking indicated by my account, without any help from sentence conflation.

Similar evidence of problems with sentence conflation comes from Sonnet 65, to which I now return. White gives a number of examples of metaphors which could plausibly have arisen by conflating two other grammatically similar sentences. However, while we are given the metaphor sentence in Sonnet 65, we are not also given the sentences conflated to produce it. Nor is this surprising: it is difficult to see how 'hold out', 'siege', 'battering' could be strung together to produce a complete sentence that comes close to matching the quite different sentence in which these words are embedded.

This problem is no doubt one of the reasons for White's talk of duck-rabbits: he wants us to think the sentence from Sonnet 65 could be read either as a description of time's decay or as a description of a siege with a battering ram. However, it is far from obvious that we could read this sentence as being one solely about sieges; unlike the Iago example, elements which are part of the 'other' situation are needed to hold the thought together.

This is not to deny that the words 'hold out', etc., make reference to some determinate situation; it is easy to imagine someone attempting to hold out against a siege mounted with, among other things, a battering ram. White's view and mine are in perfect agreement about this. But, as with (22), my view is that Sonnet 65 contains multiple links to a situation that is called on, in this case by semantic descent from 'hold out'. Once this descent is in place, the other words in the sentence make perfect sense; no need for a conflation of sentences, just a co-ordinated series of further descents to that same situation.

If things look bad for the kinds of complex metaphor White almost exclusively considers, somewhat paradoxically, it gets worse in respect of simpler cases. I have agreed, indeed applauded, White's excoriation of accounts of metaphor based almost exclusively on the use of predicate expressions in subject-predicate

metaphors. But metaphors do sometimes take this form, and the conflation of sentences seems ill-suited to handling it. For example, what pair of sentences is conflated to produce (R)?

White does acknowledge this problem, but relegates discussion of it to an appendix. His first move there is to question the importance of subject-predicate form, even claiming that concentration on it is ‘a recent phenomenon’ (White 1996: 235).¹⁸⁷ Still, even if this is historically accurate, it won't make the problems of subject-predicate metaphors go away.

White canvasses two ways of coping. The first, rather radically, suggests that theorists of subject-predicate examples, and theorists like himself, who concentrate on more complex cases, are actually writing about two different phenomena—two different figures of speech ‘which only share the name “metaphor” ’ (White 1996: 237). However, this tack is rightly abandoned in favour of a unified approach. Without complete confidence that I have understood him, that approach can be summed up this way: subject-predicate metaphors are, to use a metaphor of my own, tips of icebergs. Underlying them are fields of analogical relationships which could, and sometimes do, provide for extensions of what, in some text, is simply an ‘A is B’ metaphor. Thus, coming across: (23) Achilles is a lion, it would be a mistake to set off looking for associated commonplaces, properties generated by shift of context, or any direct substitutes for the ordinary meaning of ‘lion’. Instead, one should recognize that, even if (23) happens to be the only metaphor text, we must understand it as essentially extendable through the network of analogical relationships it depends upon. Lions behave in certain ways with each other; one could say that they display leonine fierceness, leonine courage, leonine determination, etc. In view of this, it can be said that Achilles behaves towards his comrades and enemies, as lions behave towards their conspecifics; his exercises of fierceness are human, not leonine, but they are to humankind as the leonine characteristics are to felinekind. This analogical field is perhaps only hinted in (23), though it is called upon extensively in the passages from the *Iliad* that White cites.

This way of handling subject-predicate metaphors is both surprising and disappointing. It is surprising because it seems to signal White's acceptance of things that he had been careful to reject in the body of the book. The story about analogy is a familiar one, and versions of it have been told by many Content Sufficient theorists, including those who would consider themselves heirs of Black, Beardsley, and Goodman. More important to the present discussion, however, is the fact that the proposal is disappointing: it doesn't actually justify the application of White's doctrine of sentence conflation to the case of subject-predicate metaphor. In the

¹⁸⁷ The conflation of sentences view itself appeals to an underlying duality—the similarity of situations—and this seems to be evidence of the essentially ‘this-is-that’ nature of metaphor. By itself, this doesn't justify the diet of subject-predicate examples in philosophical writing, but it certainly goes some way to explaining it as something other than a historical accident.

penultimate paragraph of the book, he finally tells us what, 'at the linguistic level', are the conflated forebears of (23). They are: (23a) Achilles is a ruthless man of war, (23b) Ferdinand is a lion.¹⁸⁸ Obviously enough it is possible to appreciate that, as far as vocabulary is concerned, (23) could be a pick-and-mix result of these two sentences. But the distance by which they fall short of sharing the motivation of the sentence conflation view is breathtaking. Is White expecting us to believe that the situations described in each of these sentences share similarities, that we are able to see one as the other? How could that be when in one case we are characterizing a man by citing one of his non-essential properties, and in the other categorizing a particular animal as a member of a certain natural kind?

Rather than pressing the failure of White's proposal, which, after all, is only offered in the book's appendix, I shall instead finish off by reminding you how my account would deal with the whole range of cases considered by White.

Semantic descent is a passage, something like reference, from words to objects. These objects can be quite various: they can be individuals, kinds, situations, actions, states of affairs, or events. What matters is not their type, but the fact that, once identified, they are pressed into service as predicates; or, more cautiously, they take over the role that predicates usually play, a role I call 'qualification'. It is the ontological differences amongst the objects of metaphor which allows the semantic descent account to deal with metaphor effects ranging from those of individual words, as in (23), to those of whole sentences. There is nothing complicated about this, simply the fact that different constituents of sentences, or whole sentences themselves, conjure up rather different objects. Despite these differences, useful as they are for giving the account flexibility, there is an underlying unity. This is the result of the fact that, while the objects differ, they all fulfil what is a single predicational function. But, as we have seen, in many cases where the syntax of the metaphor, or string of metaphors, is more complicated than that of subject-predicate, there may well be no overall assertion of the form 'A is B'. In saying what it does about summer, Sonnet 65 involves various semantic descents and qualifications, but we do not have to understand Sonnet 65 as claiming straight out that time is a battering ram. Finally, because the work begun in the words of metaphors is completed by objects, there are endless ways in which one can extend them. All you need to make sure of is that additional metaphors, either in the same sentence or in succeeding ones, descend to the original object or to another that is closely enough related to it. Failure to do so results in uninterpretable mixing of metaphors.

5.4. Indirect Speech

The last of the accounts I shall consider is Fogelin's. However, I shall not be rehashing my criticisms of his simile thesis; I have already said enough about how I think

¹⁸⁸ White 1996: 245. I thought Ferdinand was a bull, but maybe that is part of White's point. Unfortunately, this makes (23b) itself sound metaphorical.

my account accommodates the best features of this thesis, and highlights its shortcomings. Instead, I shall concentrate on a more general feature of his account, its insistence that metaphor is an indirect speech act.¹⁸⁹

Fogelin, like Davidson and Searle, thinks we should look not to the meaning of words, but rather to their use, in understanding metaphor. His belongs to that group of accounts I described in Chapter 1 as 'Content *Insufficient*'. However, within that group, he parts company early on with Davidson, throwing in his lot with Searle and others who think that we can best deal with metaphors by finding them to be, in Searle's favoured terminology, vehicles of 'speakers' meaning', and which Fogelin prefers to describe as 'indirect speech acts'. Unlike Davidson, he seems then to be an Alternative Message theorist, but, as I shall eventually argue, there are problems with this characterization.

5.4.1. Metaphor and irony

Searle, and many others tempted by the Alternative Message route for metaphor, appeal to irony as a model. Fogelin is not so tempted, and I think rightly so, but his reason for rejecting the model, and mine, are quite different. I will come to this, but first a few remarks about irony. (Do not take what follows to be an account of irony. My aim is to say just enough to make its use as a model of metaphor clear.)

If I comment on a sloppy sentence in a student's essay by saying to him:(24) You must have spent many hours constructing this sentence,I am speaking ironically. While there are lots of things that might be said about this case, this much is the minimum: any hearer who took my words at face value—who didn't recognize the need for corrective or evasive action on his part—would have missed something important, something necessary for understanding me. What is required is not a correction to the grammatical sentence I uttered, but a correction nonetheless. A hearer attuned to irony will realize that (24), taken at face value, is not an accurate representation of what I intended her to believe, or at least intended her to believe about what I believe. Some replacement sentence must be found, and taking this replacement at face value is what constitutes the needed correction. Grice himself suggests that we must replace (24) by its contradictory:(25) It is not the case that you must have spent many hours constructing this sentence,but this is less satisfactory than the more mildly contrasting:(26) You did not spend enough time constructing this sentence.

¹⁸⁹ Indirect speech act views of metaphor are often classified as 'pragmatic', and there are other pragmatic accounts around besides Searle's and Fogelin's. In particular, there is a view about metaphor that comes out of the relevance theory of Sperber and Wilson (1985/6, 1995), and it would have been good to discuss that view here. However, in order to do it justice, I would have had to take on, as background, the whole of the relevance theory, as well as some recent criticisms of that view's handling of metaphor (in Carston 2002). Given the size of these tasks, I decided to leave such discussion to another time and place.

However, for our purposes, these are mere details, and the important point is simply that, in irony, the hearer is called upon to come up with a corrective replacement that contrasts in some appropriate way with the original.

As is familiar, Searle treats metaphor as requiring pretty much the same kind of replacement strategy. Romeo says:(R) Juliet is the sun, but a hearer would be wrong to take this at face value, and is thus forced to find a replacement (or a range of replacements) which resembles (R) but which correctly represents what Romeo intends. One such replacement might be:(27) Juliet is necessary for my life. Much has been said against this account of Searle's, only a little of it in this book, and I won't even try to summarize the arguments here. However, one problem with Searle's account that was discussed briefly in Chapter 1 is of particular importance in the present context. This is its failure to meet the transparency requirement.

When we hear (or read) a straightforwardly literal sentence in a language we understand, we cannot but hear (or grasp) it as having meaning; it is in this way transparent to us. My claim was that transparency also characterizes our encounters with metaphors. Without rehashing the discussions in Chapters 1 and 3, it is worth reminding you of two important ways in which this claim was qualified. First, transparency is not an armchair substitute for psychological research into the nature and time-course of linguistic processing.¹⁹⁰ Transparency is a claim about *how* sentences strike us, not one about *when* or the *way* this happens. Second, transparency is compatible with there being more to be said about the meaning of any given metaphor. The familiar kinds of elucidation and commentary we provide for metaphors are not made redundant by the sort of understanding required for transparency.

There seem to be two ways that an account of metaphor can fall foul of the transparency requirement: by providing the wrong content, or by failing to provide any such content. The latter failing is typical of Content Sufficient accounts like Black's. Insisting that, in metaphors, familiar words have quite unfamiliar meanings hearers must somehow work out, metaphors, implausibly, become like encounters with sentences containing unfamiliar words. However, Alternative Message accounts like Searle's fall down in the other way: the contents that hearers are credited with are of the wrong sort. That they are wrong is partly because they cannot explain the kinds of thing we find it natural to go on to say about metaphors by way of elucidation or commentary. But there is in fact something more seriously amiss than this.

Adhering to the indirect speech act model provided by irony, the ultimately appropriate content of any given metaphor is the corrective one (or ones) supplied by a hearer. This corrective content is what I earlier called the 'Alternate Message', and, as the label implies, it must be an alternate to some initial content. In the

¹⁹⁰ It is all too easy to slip into inappropriately psychological talk in connection with transparency. Perhaps I am being unfair, but I find this in writers such as Recanati 2001 and Bezuidenhout 2001.

Romeo example, this initial content is generally assumed to be something like the literal claim that Juliet is the sun, that is, that she is the star at the centre of our solar system. Now, as I suggested above, I think that, if this is what we are supposed to hear in (R), then transparency has not been satisfied; such a content would not link intelligibly to the elucidations and commentaries we naturally provide for (R). However, it is possible to imagine the following response to this: given that transparency is not about nuts-and-bolts psychological processing, why not simply imagine that a hearer makes an unconscious, and almost certainly rapid, transition from initial literal content to corrective content, and that it is the latter which links intelligibly to further elucidation and commentary. At no point in any of this would a metaphor strike us as puzzling in the way that it would if it contained words we didn't understand, and that would suffice for transparency. After all, it is arguable that in irony itself, a hearer might not actually be aware of two kinds of content. Psycholinguists may one day tell us how we manage to 'get' ironic utterances, but philosophers should not construct arguments which depend upon speculations about how this is done. The plain fact is that we do make the correction necessary to the irony in (24), and do so without necessarily being aware of the stages involved in this correction. And, given its use as a model, what goes for irony, should also go for metaphor.

This response might convince some that the point about elucidation and commentary can be got around, but it exposes Searle's view to a more devastating objection, one which applies to any indirect speech act view that uses irony (or similar) as a model for metaphor. Essential to views like Searle's, and essential to the train of thought just given, is the idea that a metaphor has an initial content which is, as it were, *pro tem*, and which ultimately gives way to a corrective content. That some such content is essential is because there must be some rational transition from it to the final corrective content. Note that I am not here claiming that a hearer must carry out the reasoning that figures in this transition, or even that the hearer must be aware that it could be done. These are psychological matters that can be left on one side. Instead, what is at issue is simply the fact that a corrective content is a *correction*: unless there is something that is corrected, the final content of the indirect speech act simply wouldn't be a correction. Further, and crucially, the ground for moving from the original to a corrective content must be rationally constructible (even if a hearer doesn't construct it himself). This is all but explicit in the requirement that the final content is a *better* option for capturing a speaker's intentions.

Just for the purposes of contrast, think how different it is for a certain kind of Content Sufficient theorist. On his picture, the materials from which a hearer constructs a metaphorical meaning are the familiar words that typically figure in metaphors. In ordinary contexts we know what kind of contribution they would make to the content of speech acts in which they figure. But in a metaphor, we can no longer rely on these familiar meanings, and our task is that of providing new, metaphorical meanings for the relevant words (which stripped of ordinary meaning are in fact more like word-forms). If we succeed in this, we then do have a content for the speech act made by uttering the metaphor sentence. But this content need

not be the result of some process of reasoning from the materials we begin with: brute associations are more likely to play a role than argumentative transitions.

The difference between the commitments of these two kinds of account is crucial. While some Content Sufficient theorists, and even a No Message theorist like Davidson, might allow that there really is a content in the words of a metaphor taken literally, even if it is hopeless for understanding the metaphor, there is nothing about Content Sufficiency that *requires* this. But there is no such leeway in respect of Alternative Message accounts like Searle's. Were there to be no original content, standing in need of correction, we would have to reject any such view out of hand. Nor is this a mere possibility: if we put on one side the simplest subject-predicate metaphors, there is simply no reason to think that metaphors have anything worth calling literal or initial content.

This point is made by White (1996: 204–5f), though I do not owe it to him.¹⁹¹ Among the metaphors he uses to make the point is: (28) They ought to donate his face to the wildlife fund.¹⁹² We might think that 'Juliet is the sun' might be taken to be contentful but literally false, but it doesn't seem possible to take (28) in a similar way; if one leaves metaphor behind, (28) is sheer nonsense. This same fate awaits virtually any of my examples in Chapter 4. I think that an honest look at each of: (29) Out of the crooked timber from which men are made, nothing straight can ever be built. (30) In cities you build a language of circumspection and tact, a thousand little intimations, the nuance that has the shimmer of rubbed bronze. (31) Swerving at the last moment to avoid innocent bystanders, his argument came to halt. (32) Her prose shows traces of the rough timbers that a more careful builder would have covered over. would reveal no way in which each could be taken to be literally false. It is not even clear what it would mean to take them literally, or to take them as conveying any kind of non-metaphor content.

White suggests that theorists whose views are challenged by these examples could have a fall-back position. Instead of metaphors being counted literally false, and therefore in need of some kind of replacement, why not count their being *either false or nonsensical* as sufficient motivation for further interpretative work? Aside from any desperation evident in this, it just won't work for the kind of Alternative Message account I am presently considering. White notes that Searle allows himself to speak at one point of the literal readings of metaphors as 'semantic nonsense', but of course if they were really nonsensical, we couldn't reconstruct the reasoning

¹⁹¹ White cites Levin 1988 as the only other author he knows of to have made this objection to Davidson.

¹⁹² Said by Mohammed Ali about Joe Frazier before their fight in Africa. See White 1996: 205.

which leads from them to the proper corrective content. It is essential to irony that there be some reading of the ironic sentence which can serve as input to an argument whose output is the appropriate speaker meaning. Whether that input is best thought of as contradicted, or made 'opposite' in some other way, by that argument I leave to connoisseurs of irony, but if the input is nonsense, then it could not serve in the premises of any such argument. A view like Searle's requires the words of a metaphor to make some kind of sense even if it is not yet the right kind. But this requirement is rarely fulfilled.¹⁹³

5.4.2. Metaphor without irony

I have been careful to insist that the above objection applies only to Alternative Message accounts which assimilate metaphor to figures like irony. Fogelin resists any such assimilation, and his view thus escapes my objection. However, there is something of the frying pan and the fire about this escape.

The key feature of the irony model that causes trouble is the idea that the corrective action on the part of hearers requires supplying a replacement for a content already grasped. But, while aware that some correction must be present for metaphor to count as an indirect speech act, Fogelin is very careful to deny that, in metaphor, the correction involves a *replacement* content. In a key passage, he writes: What we want may, at first sight, seem impossible to get: a theory that allows us to say that an utterance when taken literally is false, but when taken metaphorically is true, even though there has been no shift in the meaning in these two ways of taking the utterance. In fact, however, the traditional elliptical-simile theory solves this dilemma in a straightforward and natural way. A metaphorical utterance of the form 'A is Φ ' just means, and literally means, that A is like Φ . Likeness claims, however, have criteria of adequacy that shift with context. If someone says that A is like B in one context, and then says it again in another, then although he has said the same thing twice over (that A is like B), one of these utterances could be true while the other false. If we like, we can still talk about a shift taking place, but it is not a shift in the meaning of words; it is a shift ... in the modes of relevance and evaluation governing the likeness claim. (Fogelin 1988: 75–6) Instead of the replacement we have in irony, the correction in metaphor consists in a shift of context. But there is something puzzling about this, something suspiciously like a sleight of hand.

Suppose, first, we realize that the proper context to apply to a particular instance of an utterance 'A is Φ ' is one which will yield the fully metaphorical meaning. Then, given that we go straight to the interpretation of the utterance as a metaphor, why is this an indirect speech act? We hear the words, access the proper context, and thereby generate an interpretation. This sounds very much like a case of a context-sensitive *direct* speech act. There is, in short, no correction made here.

Alright, let's suppose instead that we start off by interpreting 'A is Φ ' in a context which is appropriate to a *literal* rendition. As we have seen, this leads either to a

¹⁹³ The objection here is stronger than one usually made to Searle's view. Whereas it is noted that the reasoning which leads from initial to corrective content in metaphor lacks the determinacy of the reasoning in the case of irony, I am claiming that in metaphor there is often nothing to reason *from*.

falsehood, or, more likely, to nothing at all sensible, and perhaps that is why we find ourselves hastily shifting to a different context in continuing our interpretative efforts. But now we have a problem. If the initial literal interpretation is false, then the correction involves a replacement. We have an implausibly false sentence when we wouldn't have expected it, and we change the context of interpretation so as to come up with a reinterpretation that is more plausible, perhaps even true. However, this couldn't be what Fogelin has in mind in the passage above, since what I have just described is tantamount to Searle's irony view. What about the following alternative? We set out to interpret in the literal context, but get nowhere; we get no interpretation of the words in the utterance. As White and I agree, this is the more realistic alternative. Faced with this impasse—this inability to interpret the utterance at all—we see it as a reason to switch to a different context, and this switch is rewarded by an intelligible metaphorical interpretation. This sounds more like what the passage above requires, but it is really no improvement on the very first alternative described above—the one in which we simply set out to interpret by using the non-literal context in the first place. To be sure, unlike that case, we switch contexts because we have failed to find any sensible interpretation using the literal one, whereas in the earlier alternative, I never said why we went straightaway for the non-literal context. But for all practical purposes they are the same: we hear an utterance which we can only interpret using a non-literal context. However, this means, as it did above, that, because there is no *direct* speech act, the metaphorical reading cannot be *indirect*. (We hear the words *uttered*—there is, if you will, a direct *utterance* act—but if we took this seriously, we would end up saying that every speech act was indirect.)

The bottom line of all this is actually rather surprising, though also rather satisfying. Fogelin's view, for all that he insists that metaphor is an indirect speech act, is in the end rather more like Stern's view than it is like Searle's. This is sort of signalled when, some pages after the passage above, Fogelin writes: I asked how it could be possible for an utterance to be false when taken literally, but true when taken metaphorically without there being any shift of meaning from the one reading to the other. I can first note that parallel situations arise with non-figurative language ... [C]loser to the present case, if 'x is good' means something like 'x satisfies relevant standards of evaluation', then saying x is good could be true in some contexts, but not others, without there being any shift in the meaning of what is said. (Fogelin 1988: 91)¹⁹⁴ There are problems in this passage with Fogelin's talk of 'what is said' and, while approving of the appeal to context-sensitivity, Stern is quick to point them out. But I think the way in which Fogelin's view shadows Stern's transcends these problems, and is deeper than Stern realizes. Stern's own view required him to invent an operator, 'Mthat' which, in order to deal with metaphor, imposes context-sensitivity on words not ordinarily context-sensitive. Fogelin insists that metaphors are implicitly similes: a metaphor 'A is Φ ' is equivalent to the simile 'A is like Φ '. He also insists

¹⁹⁴ Stern 2000: 219 cites this same passage.

that we can only understand a simile correctly if we recognize that 'like', in the simile, forces a change of context on us. There are thus grounds for thinking that Fogelin's figurative 'like' and Stern's 'Mthat' do something like the same job. Each is called into action when an utterance is taken to be metaphorical, each marks a change in sensitivity to context in comparison with wholly literal utterances. Of course, there are real differences, and, by and large, I think that Stern has a much clearer view of the territory. But even if the similarity is only fleeting, it is relevant that Stern insists his view is, as he puts it, a semantical one; in my terms, it is Content Sufficient. For, as I argued above, though Fogelin sets out his view in terms of indirect speech acts, metaphors end up sounding like direct speech acts in which context-shifts play a crucial role. In my terms, Fogelin sees himself as offering an Alternative Message account, but the result is more Content Sufficient than anything else.

Epilogue: Language and Thought

From the time that its outline first took shape, I thought of my view of metaphor as obvious,¹⁹⁵ and the puzzle was to understand why it hadn't occurred to anyone before. In Chapter 5, I considered three recent views which contain elements of my account, and overlap with it. On delving into them, I remember being convinced at every turn of the page, either that my view would be there in black and white, or that definitive grounds for rejecting it would emerge. Fortunately, neither of these alternatives seems to have been realized, but this still left the puzzle of why this should be so. Nowhere was this more forcefully brought home than by Henle's well-known article on metaphor. I had read it a long time ago, and it is often cited for lots of different purposes in the literature on metaphor. But given that I had, by the time of my rereading, come to think along the lines of the semantic descent account, imagine my reaction to reading: Metaphor, then, is analyzable into a double sort of semantic relationship. First, using symbols in Peirce's sense, directions are given for finding an object or situation. This use of language is quite ordinary. Second, it is implied that any object or situation fitting the direction may serve as an icon of what one wishes to describe. The icon is never actually presented; rather, through the rule, one understands what it must be, and through this understanding, what it signifies. (Henle 1958: 178) Here we have a view of metaphor as involving, first, the obtaining of some object by the very process I call semantic descent and, next, a use of that object which is described as 'iconic' and also 'semantic'. Is this not my view?

While it certainly gave me pause, and whilst I do think that Henle's view sounds startlingly close to mine, I think his view is not mine. Saying why this is so will go some way towards explaining why the semantic descent account hasn't figured in the literature before now, and will at the same time prepare the way for the central point of this Epilogue.

Henle spells out what is involved in taking one object as an icon of another in terms of similarity, and in doing so he cites Peirce's work on signs. Peirce says of an icon that 'it represents its object in resembling it' (Peirce 1966: 368), and Henle writes: a metaphor, as distinguished from the other tropes, depends on analogy, and in this analogy one side is used to present the other. Thus envelopment in a cloak is used to present the notion of gloom [metaphor: 'And hateful thoughts enwrap my soul in gloom']; the character

¹⁹⁵ This is not to say that my account is correct. The issue here is its obviousness as a *possible* account, not the obviousness of its truth. Given the effort I have put into defending it, you will recognize that I find it true, but not obviously so.

of a man is presented through its likeness to a fox; and the appearance of a worm through its likeness to a bit of rubber [metaphor: 'an obliging thrush hopped across the lawn; a coil of pinkish rubber twisted in its beak']. In each case we are led to think of something by a consideration of something like it, and this is what constitutes the iconic mode of signifying. (Henle 1958: 177)¹⁹⁶In being struck by the overlap between Henle's first stage and my semantic descent, it is all too easy to gloss over the differences at the second stage. But these passages should make absolutely clear not only the gap that exists between Henle's notion of an icon and my notion of qualification, but also the inadequacy and confusion that attends Henle's use of Peirce's notion in connection with metaphor.

An icon is a way of moving from one object to another, it is first and foremost a referential notion. Indeed, all three of Peirce's kinds of sign—indices, icons, and symbols—are referential, differing only in how one is taken from sign to referent. Indices refer directly via some 'real' relation, icons via similarity, and symbols depend upon some established conventional linkage. Peirce belongs to the army of philosophers, and non-philosophers, who think of the language–world relationship in terms of reference. It might seem that because icons are signs which reach out to their referents via something like the information they contain, we have here the first step on the road to predicate independence. But this is illusory: in the end it is resemblance—the putative sharing of properties—which links iconic objects and their referents, rather than anything like the information that these objects might bring to bear.

Henle denies that metaphor is implicit, elliptical, or condensed simile, but he, like so many others, thinks that we can best explain metaphor in terms of a sharing of properties. For him, the object accessed in metaphor—the one got by semantic descent—shares properties with the metaphor subject, and it is because of this that the first 'presents' the second. He never notices how odd it is to say that, for example, a cloak shares properties with gloom, and his gloss on this metaphor is instructive: Gloom is a pervasive affair in that it influences one's entire mental outlook. It need not have consciously felt causes or itself be the object of direct awareness but rather it tinges all other mental activities. In these respects it is like a wrapping which covers the whole of an object and which allows its form to show through, though modified by the covering. (Henle 1958: 180) Gloom is like a cloak? Perhaps, but only if one has already taken 'like' in a figurative sense, that is, a sense that can only be understood after the processes of metaphor have done their work in transforming the properties of a cloak and the properties of gloom, so that they can be said to be shared. Moreover, even if one does end up finding there to be certain metaphorically transformed properties shared by gloom and cloaks, the function of metaphor remains a puzzle. Surely, there is more to Keats's metaphor than using one object to 'present' another. (Synecdoche and metonymy)

¹⁹⁶ Henle's cloak metaphor comes from Keats's 'To Hope', and the worm metaphor is from Woolf's *Between the Acts*.

might fit in with talk of presentation and, as I allowed, they belong with metaphor. But the cloak/gloom case seems to be a case of metaphor proper.)

Where Henle thinks of the object got by semantic descent as an icon, as referring to the metaphor subject in virtue of sharing properties with it, I see the object as a non-linguistic predicate. Keats enlists the situation of being enwrapped in a cloak as a way of informing us about gloom. In so informing us, this cloak-enwrapment takes on the role usually played by the form of words we call a 'predicate'. To avoid confusion, I was forced to invent a term for this bit of extra-linguistic informativeness, namely 'qualification'. Of course, once we see things this way, it is easy to understand why talk of property-sharing is tempting: if cloak-enwrapment is, as I think we should say, true of gloom, then there is bound to be something we can also say about the properties of one which are shared by the other. Indeed, the bit I quoted above from Henle is a perfect example of at least part of the kind of elucidatory commentary that I locate downstream of metaphor. In such a commentary, one makes remarks about the properties of a metaphor-object with a view to encouraging the figurative transformation of them into properties shared by the metaphor-subject. However, in being downstream—in depending on the metaphor already being in place—the shared properties that result from this transformation cannot themselves account for the metaphor. (It should also be remembered that talk of property-sharing in respect of metaphor is not always appropriate. Think here of the syntactically more complex cases, discussed in Chapter 4, in which events and states of affairs are metaphor vehicles.)

The role that qualification plays in helping us understand metaphor has been central to this book, and it is the one feature of my account which seems not to have been anticipated in the literature. But there are good reasons for this, reasons which serve to explain why writers veer away from the full semantic descent account, even when, as in Henle, it seems to be right there.¹⁹⁷

In Chapter 2, I cited the rhetoric deployed in philosophical logic in favour of treating reference and predication as making separate and equal contributions to the primitive unit of language and thought which Quine and Strawson call the basic combination. But I also noted that equality of treatment rarely goes beyond rhetoric when it comes to predication. Linguistic predication is itself mostly explained in terms of reference (the *specification* of concepts by predicate expressions is crucial to it) and, whereas we think of reference as a function that can be fulfilled by non-linguistic objects, one is likely to be puzzled by the very idea of non-linguistic predication. As a result of all this, my missions in Chapter 2 were, first, liberating predication from the yoke of reference and, second, revealing the ways in which non-linguistic predication figures in our practices, even if we don't ordinarily think of it that way, or call it that.

¹⁹⁷ Alston 1964: 98–9 bases his view of metaphor entirely on Henle's. But in restating the view, he gets even closer to my idea of qualification than Henle. He writes of Shakespeare's 'sleep that knits up the ravelled sleeve of care': 'by considering this situation we can realize something of the effect of sleep on a careworn person', and then he goes on: 'But this is only possible if there is some important and readily noticeable similarity between these two situations'. It would be difficult to find more direct evidence for the 'veering' mentioned in the text.

Given the obstacles to seeing predication as radically distinct from, but wholly equal to, predication, it is not in the end surprising that the semantic descent account has not figured up to now in the literature on metaphor. Every other aspect of my account has figured, explicitly or implicitly, but just when the idea of qualification comes over the horizon, prejudices perhaps inherited from philosophical logic make it all but invisible.

From the very beginning, I insisted that the phenomenon of qualification is not itself metaphor, nor does it depend on metaphor for either its existence or importance. However, being now at the end of the book looking beyond metaphor, it is necessary to return to the notion of qualification, and to leave metaphor behind. My contention is that qualification has substantial consequences for philosophy (and psychology), though in this Epilogue I can only hint at what these might be. I shall begin by making a few remarks about a certain interdependence thesis—the interdependence of reference and predication. (This interdependence thesis figured only briefly in my efforts at predicate liberation in Chapter 2.)

The rhetoric that stresses the joint contribution of reference and predication to the basic combination certainly suggests that these functions are not only dependent, but *interdependent*. The headline view is that, at the most primitive level of language which is the basic combination, we find both reference and predication, each of which is necessary for the basic combination, and each of which depends for its functioning in the basic combination on the other. Given this, it would seem plausible to insist that reference and predication are interdependent; that one couldn't so much as understand reference without predication, and vice versa. But, as we saw in Chapter 2, appearances can be deceptive. One's attitude towards the interdependence thesis is apt to be shaped by the prejudices one has in respect of predication. If you think of a predicate as a form of words in a natural language, but think of reference as something which can be achieved both inside and outside natural language, and if you also think of reference as underpinning natural language predicates, then you are likely to find the interdependence thesis trivially true, and hence not all that useful.¹⁹⁸ However, if your consciousness was raised in Chapter 2 in respect of predication, and you are prepared to see it as an equal of reference, then the interdependence thesis comes into its own. Far from being a trivial consequence of reference's role in our understanding of predicate expressions, it announces a substantial constraint on what is required for the possession of linguistic abilities. It insists that we cannot so much as recognize an ability as one of

¹⁹⁸ Davidson 1984*b* insists on the primacy of truths, and this suggests a way to argue for the interdependency thesis: if both reference and predication can only be found by a kind of subtraction from whole truths (basic combinations, among other things), then obviously one couldn't have reference without predication or vice versa. However, from my point of view this approach is problematic. I have argued for, and continue to think it important, to accord to predication the same rights and privileges as are accorded to reference. Though he doesn't quite put it this way, Davidson's view assumes that even reference hasn't rights and privileges in the first place; equality thus comes about by the denial of rights to reference rather than by their provision to predication.

referencing, unless we can also find evidence of an independent and correlative ability to predicatively characterize, whether such characterization is accomplished with natural language predicates or by using objects.¹⁹⁹

So understood, I think that the interdependence thesis can give us considerable leverage on certain questions about language and thought. I have in mind here: the question of whether there can be animal languages; the question of how certain languageless human beings (infants) can come to possess language; and finally the question of whether there can be thought without language. In one way or other, each of these questions is about boundaries. Where and how do we draw a line between creatures with, and those without, language? What makes it possible for a human infant to cross this line? Should crossing this line also be a requirement for the possibility of thought? If the interdependence thesis is interpreted in the way I favour, and if it is true, it should be more difficult to cross these sorts of boundary, and also easier to tell if any attempt at crossing succeeds. (These are both important 'ifs'.)

As an illustration of the uses and misuses of interdependence, consider the first of the boundary questions—the one about animal language. Suppose you are wondering, as many have, whether there could be a kind of animal, less sophisticated than *homo sapiens*, which might nonetheless have a language. Many think that chimpanzees are actually such a kind. Using them as a test case, suppose you find that, after training, these creatures learn how to manipulate various plastic tokens in ways that lead us to think of these tokens as sort of like labels, and therefore as referential. After this training, a chimp is perfectly capable of employing particular tokens in circumstances that seem to show the token to stand for, for example, some piece of food, some particular lab assistant, some location, etc. Are these tokens genuinely names of these things? Is this acquired ability an ingredient or element of a language? Much ink has been spilt on these questions, so the following small additional contribution should seem negligible.

If you think that a pattern of behaviour can only count as referring if it takes place alongside an additional pattern understood as predicative—and vice versa—then you may well be unconvinced by the description of the chimps' manipulation of tokens. Where, you will ask, is there any evidence of predication? All you can see is an ability to connect tokens and particulars, and, even though the chimps are very good at it, these connections will not by themselves constitute evidence of language. They might well just be causal associations. Remember, the minimal condition for language is facility with the basic combination, and the minimum necessary for this is a facility with both reference and predication. If the chimps engage in something which resembles reference, but if they do not ever use this in producing basic combinations—something that would also require predication—then, by the interdependence thesis, we should be sceptical in regarding the linkage between plastic tokens and particulars as referential after all.

¹⁹⁹ Vice versa too, since the thesis is one of *interdependence*. That is, our characterizing ability can only be discerned against the background of our ability to refer. I relegate this to a footnote because I cannot imagine anyone challenging the thesis in this direction.

As so deployed, the truth of the interdependence thesis promises help in settling this boundary issue, but, as will be familiar from the literature, matters are not so simple. It is just at this point that the usual misunderstanding of predication surfaces, and messes up this application of the interdependence thesis. On the one hand, it will be said that predication is something essentially linguistic; that in requiring evidence of this in chimps, we are begging the question of whether their apparently referential activity is evidence of any kind of linguistic ability, however primitive. On the other hand, there is the confounding influence of the near universal tendency to treat predication in terms of some kind of referential relationship. This second point needs to be spelt out.

Above, I had assumed that the chimps linked their plastic tokens with particulars. But what if one thought of these tokens as more like concepts, and also thought that predication could be understood in terms of concept specification? A chimp uses a particular token to bring our attention to a piece of fruit (or so it seems). If this token is understood as, say, 'banana', then the way is open for us to think of its use as something like that in the one word sentence, 'Banana!' Familiarly, this latter construction is often understood as: give me this thing which is a banana. And if it is so understood, it is not difficult to imagine someone insisting that the chimp thereby shows a facility with both predication and reference, and therefore a fully, if primitive, linguistic facility with what appears to be a basic combination. The readiness of psychologists (and philosophers who comment on these cases) to treat the tokens as involving both reference to particulars and reference to concepts makes it impossible to use the interdependence thesis in any substantial way. In effect, the thesis reverts to its trivialized version. In order to restore the substantial and useful version, we must desist from treating these tokens as uses of predicates like 'banana', merely because they apparently refer to things which we know to be bananas; we have to find a way to discern predication which does not depend simply on reference. In short, we must take seriously the difference, independence, and equality of the linguistic role that the rhetoric of philosophical logic accords to predication in the first place.

To repeat, what is required to give substance to the interdependence thesis is the possibility of there being some activity we can recognize as predicative, even if in some primitive form, and which is independent both of natural language and of the influence of reference on the very idea of predication. It is here that my notion of qualification can help, though it will not be straightforward. What is needed is some story about the evidence we might uncover which would show that a creature—by hypothesis, one without any natural language—in fact uses objects other than words to effect both acts of reference and acts of qualification. Given the interdependence thesis in its substantial form, this evidence would most certainly not be separately discoverable in respect of each of these activities, and it is admittedly not easy to imagine what it would be like. What kind of performance on the part of the chimps would we take as showing that they could be credited with a certain way of picking out particulars (reference), but also with a way of characterizing them that did not itself call on the resources of natural language?

Difficult as this first boundary question is, questions and their difficulties multiply when we come to think about the other boundaries. We might be prepared to say that chimps don't actually have a language, precisely because there is nothing they could do which would be evidence of genuine predicative abilities. This doesn't seem startling because there are many other reasons to be sceptical about the claims for chimpanzee language. But we have to be careful to insulate this response from the other questions. Whatever else is true, children come to be natural language users, even if, from some points of view, they start out little ahead of chimps. And, if a head count means anything, there are many more philosophers who insist that there can be thought without language than there are those who find language necessary to thought, and these former will not be as impressed as I think they should be by the interdependency thesis. Still, even though these questions are more difficult than the one about animal language, I believe that the interdependency thesis, and the role of qualification in making that thesis substantial, can help. At the end of Chapter 4, I suggested that qualification could go some way to justifying the thought that metaphor is actually rather basic to language generally. The point of this Epilogue has been to suggest that this same notion might help us shift some of these boundary questions about language, and thought and language, out of the ruts they have made in philosophical discussion in recent years.

References

- Alston, W. P. (1964) *Philosophy of Language*. Englewood Cliffs, NJ: Prentice Hall.
- Armstrong, D. M. (1997) *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Beardsley, M. (1962) 'The Metaphorical Twist', *Philosophy and Phenomenological Research*, 22, 293–307.
- Bezuidenhout, A. (2001) 'Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor', *Midwest Studies in Philosophy*, 25, 156–86.
- Black, M. (1954–5) 'Metaphor', *Proceedings of the Aristotelian Society*, 55, 273–94. (Reprinted in Black 1962.)
- (1962) *Models and Metaphor*. Ithaca, NY: Cornell University Press.
- (1979/1993) 'More about Metaphor', in A. Ortony (ed.), *Metaphor and Thought*. Cambridge: Cambridge University Press.
- Blackburn, S. (1984) *Spreading the Word: Groundings in the Philosophy of Language*. Oxford: Oxford University Press.
- Bolinger, D. (1980) *Language: The Loaded Weapon*. London: Longmans.
- Carston, R. (2002) *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell Publishing.
- Cohen, L. J. (1979/1993) 'The Semantics of Metaphor', in A. Ortony (ed.), *Metaphor and Thought*. Cambridge: Cambridge University Press.
- Cooper, D. (1986) *Metaphor*. Oxford: Blackwell Publishing.
- Davidson, D. (1984a) 'What Metaphors Mean', in his *Truth and Interpretation*. Oxford: Oxford University Press.
- (1984b) 'Reality without Reference', in his *Truth and Interpretation*. Oxford: Oxford University Press.
- Davies, M. (1982/3) 'Idiom and Metaphor', *Proceedings of the Aristotelian Society*, 83, 67–85.
- Delillo, D. (1999) *Underworld*. London: Picador.
- Fogelin, R. (1988) *Figuratively Speaking*. New Haven: Yale University Press.
- Frege, G. (1952) *Philosophical Writings*, ed. Max Black and Peter Geach. New York: Philosophical Library.
- Gibbs, R. (1994) *The Poetics of Mind*. Cambridge: Cambridge University Press.
- Kushner, J., and Mills, R. (1991) 'Authorial Intentions and Metaphor Comprehension', *Journal of Psycholinguistic Research*, 20, 11–30.
- Goodman, N. (1972) 'Seven Strictures on Similarity', in his *Problems and Prospects*. Indianapolis, Ill: Hackett Publishing.
- (1976) *The Languages of Art*. Indianapolis, Ill: Hackett Publishing Co.
- (1984) *Of Mind and Other Matters*. Cambridge, Mass.: Harvard University Press.
- Grice, H. P. (1957) 'Meaning', *Philosophical Review*, 66, 377–88.
- Guttenplan, S. (1992) 'Wittgenstein contra Chomsky', in D. Andler, P. Jacob, J. Proust, F. Recanati, and D. Sperber (eds.), *Épistémologie et Cognition*, Liège: Mardaga.
- Hale, R., and Wright, C. (1997) *Companion to the Philosophy of Language*. Oxford: Blackwell Publishing.
- Henle, P. (1958) 'Metaphor', in P. Henle (ed.), *Language, Thought and Culture*. Ann Arbor, Michigan: University of Michigan Press.

- Kant, I. (1784/1912) 'Idee zu einer allgemeinen Geschichte in weltbürglicher Absicht', *Kant's gesammelte Schriften*, viii. Berlin: Preussische Akademie.
- Kittay, E. (1987) *Metaphor: Its Cognitive Force and Linguistic Structure*. Oxford: Oxford University Press.
- Lakoff, G., and Johnson, M. (1980) *Metaphors we Live by*. Chicago: University of Chicago Press.
- and Turner, M. (1989) *More than Cool Reason: A Field Guide to Poetic Metaphor*. Chicago: University of Chicago Press.
- Levin, S. (1988) *Metaphoric Worlds*. New Haven: Yale University Press.
- Lewis, D. (1994) 'Reduction in Mind', in S. Guttenplan (ed.), *Companion to the Philosophy of Mind*. Oxford: Blackwell Publishing.
- McDowell, J. (1998) *Mind, Value and Reality*. Cambridge, Mass.: Harvard University Press.
- Moran, R. (1997) 'Metaphor', in Hale and Wright 1997.
- Nabokov, V. (1966) *Speak Memory*. London: Everyman.
- Ortony, A. (1979/1993) *Metaphor and Thought*. Cambridge: Cambridge University Press.
- Peirce, C. S. (1966) *Selected Writings: Values in a Universe of Chance*, ed. P. P. Wiener. New York: Dover Books.
- Quine, W. V. O. (1970) *The Philosophy of Logic*. New York: Prentice Hall.
- Ramachandran, V. S., and Hubbard, E. (2003) 'Hearing Colours, Tasting Shapes', *Scientific American*, 288, 43–9.
- Recanati, F. (2001) 'Literal/Non-Literal', *Midwest Studies in Philosophy*, 25, 264–74.
- Rich, A. N., and Mattingly, J. B. (2002). 'Anomalous Perception in Synaesthesia: A Cognitive Neuroscience Perspective', *Nature (Reviews)*, 5, 43–52.
- Roth, P. (2001) *The Human Stain*. New York: Vintage Books.
- Searle, J. (1979/1993) 'Metaphor', in A. Ortony (ed.), *Metaphor and Thought*. Cambridge: Cambridge University Press.
- Shakespeare, W. (1988) *The Oxford Shakespeare: Complete Works*, ed. S. Wells and G. Taylor. Oxford: Oxford University Press.
- Sperber, D., and Wilson, D. (1985/6) 'Loose Talk', *Proceedings of the Aristotelian Society*, 86, 153–71.
- and — (1995) *Relevance: Communication and Cognition*, 2nd edn. Oxford: Blackwell Publishing.
- Stern, J. (1985) 'Metaphor as Demonstrative', *Journal of Philosophy*, 82, 677–710.
- (1991) 'What Metaphors do Not Mean', *Midwest Studies in Philosophy*, 16, 13–52.
- (2000) *Metaphor in Context*. Cambridge, Mass.: MIT Press.
- Strawson, P. F. (1974) *Subject and Predicate in Logic and Grammar*. London: Methuen.
- Suhamy, H. (1981) *Les Figures de style*. Paris: Presses Universitaires de France.
- Swift, J. (1726/1985) *Gulliver's Travels*, ed. Peter Dixon and John Chalker. London: Penguin.
- Thomas, Dylan (2000) *Collected Poems 1932–1953*. London: Phoenix.
- Travis, C. (1981) *The True and the False: The Domain of the Pragmatic*. Amsterdam: John Benjamins.
- Urdike, J. (2002) *Seek my Face*. London: Penguin.
- White, Roger (1996) *The Structure of Metaphor: The Way the Language of Metaphor Works*. Oxford: Blackwell Publishing.
- Wiggins, D. (1984) 'The Sense and Reference of Predicates: A Running Repair to Frege's Doctrine and a Plea for the Copula', *Philosophical Quarterly*, 34, 311–28.
- Wittgenstein, L. (1914–16/1969) *Notebook*, tr. G. E. M. Anscombe. Oxford: Blackwell Publishing.
- (1922/1961). *Tractatus-Logico-Philosophicus*, tr. D. Pears and B. McGuinness. Oxford: Blackwell Publishing.
- (1953) *Philosophical Investigations*, tr. G. E. M. Anscombe. London: Routledge & Kegan Paul.

This page intentionally left blank

Index

- adjective-noun and metaphor, 173–5
Ali, Mohammed, 287
allegory, 5
Alston, W. P., 293 n. 3
Alternative Message account, 9, 248, 287; *see also* Indirect
 Speech account
ambiguity, 140
animal language, 89–90, 295–7
Aristotle, 159, 218
Armstrong, D., 48
assertion and metaphor, 15–16, 121–9
attunement, 124–5, 135–6, 143–56, 178; and context,
 146–56; and meaning, 143–5
basic combination, 43–59
Beardsley, M., 13, 81 n. 36, 121, 157, 248, 282
Bezuidenhout, A., 285 n. 18
bifurcation (and oscillation), 166–70, 176
bipolarity, 159–65
bisemy, 140
Black, M., 13–14, 20, 28–30, 121, 149, 153, 157, 183–4, 248,
 257–8, 267, 282, 285
Blackburn, S., 5, 242, 244–6, 249 n. 1
blue moods, 237–8
Bolinger, D., 242–3
brute similarity, 270–3
Carroll, L., 197
Carston, R., 284 n. 17
Casablanca, 218, 225–6
catachresis, 5, 221
categorization, 78–82, 200–2; and qualification, 78–82, 84,
 105, 200–2; and similarity, 80–1
character, 250
Chomsky, N., 144, 146, 148
Cognitive account, 187–91
Cohen, L. J., 14
comparisons and simile, 203–17
compositionality, 148–9
concept exemplification, 44
concept specification, 43–4
concepts as creatures of language, 49, 51, 53, 88–91
concept-words, 43
conceptual art, 231
Conflated Sentence account, 266–83
Content Insufficiency, 8–9, 12–13; and dead metaphor,
 184–5
Content Sufficiency, 8–9, 13–14, 121, 248, 257, 286–7; and
 context, 150–1; and dead metaphor, 183–5
context, 63–4, 119–20, 148–56, 249–50, 252–5, 263–6; and
 attunement, 145–56
Cooper, D., 178, 185–8, 193 n. 24, 223 n. 44, 242–3, 249 n.
 1
culture, 147; and training, 146–8
Cummings, L., 218
Davidson, D., 3 n. 3, 9, 12–13, 15–20, 27 n. 20, 28, 62 n. 18,
 95–6, 104, 112 n. 12, 141, 151, 160, 183–5, 248, 249 n. 1,
 267, 274, 284, 287, 294 n. 4
Davies, M., 13 n. 5, 17
Davies, P., 218
dead metaphor, 154, 182–202, 221, 240; and the Cognitive
 account, 187–8; and mixed metaphor, 185–6; and
 semantic descent, 191–202
Delillo, D., 21 n. 15, 158 n. 1
Demonstrative account, 249–66; and the Semantic
 Descent account, 260–3, 265–6
demonstratives, 65, 70, 152 n. 40; and descriptions, 251–2,
 255–6; and metaphor, 96–105; and properties, 102–5
descriptions and exemplars, 105–12
Douglas, M., 147
down-arrow notation, 117; and quotation, 117, 261–3
Dthat, 251–2, 254–6
duck-rabbit and metaphor, 269
Duncan, I., 218, 224–6
elucidation of metaphor, 135–6; and extended metaphor,
 182
ex metaphor, 192–3, 222

- exemplification, 39–42, 91–2, 225–6, 258–9; and reference, 64–5
 extended metaphor, 181–2; and the Cognitive account, 191
 family resemblance, 123
 fat Wednesday (and lean Tuesday), 238–9
 fictional objects, 115
 figurative, 5, 221–2, 231–9; and literal comparisons, 204–6; and the non-literal, 210–11, 231–2
 Fodor, J., 123–8
 Fogelin, R., 33 n. 25, 81 n. 36, 189 n. 22, 203–5, 210 n. 35, 211 n. 37, 270–4, 279, 283–4, 288–90
 Frazier, J., 287
 Frege, G., 44–51, 53, 85, 107 n. 9
 French, P., 218
 general concepts, 43
 Gibbs, R., 174 n. 12
 Glancey, J., 69 n. 27
 Goodman, N., 14, 39–42, 62, 64–8, 82, 91–2, 121, 229–30, 257–9, 282
 Grice, H. P., 13, 15, 61 n. 17, 71
 Guttenplan, S., 146 n. 37
 Hayden Planetarium, 69
 Henle, P., 291–3
 Hobbes, T., 242–3
 Hollywood, 218, 225
Hours (The), 218, 229
 Hubbard, E., 233–6
 hybrid predicates, 99, 122, 126–8
 hyperbole, 5, 221, 223
 icons, 292
 idiom, 196, 222, 240
Iliad, 268, 273, 276
 Image account, 26–8; and the Semantic Descent account, 138–43; and the Transparency requirement, 138–43
 Indexicals, 250–2
 Indirect Speech account, 283–90
 Interdependency Thesis, 294–7
 irony, 5, 221, 223
 Johnson, M., 186–90, 202
 Kant, I., 21, 158, 178
 Kaplan, D., 250–7
 Keats, J., 292–3
 Kittay, E., 14, 33 n. 24, 121
 Lakoff, G., 186–90, 202
 Levin, S., 287 n. 19
 Lewis, D., 91 n. 43
 ‘like’, 203–17; as a hedge, 170–3, 206–8; literal vs figurative, 210–11; and the Semantic Descent account, 207–17
 linguistic control, 155, 161; of qualification, 99–101
 literal and figurative comparisons, 204–6
 literal vs metaphorical, 30–3
 litotes, 5, 221
 Locke, J., 74
 McDowell, J., 124, 128
 Mattingly, J. B., 236
 meaning, 7–9; and assertion, 124–8; and attunement, 143–5; as extension-determining, 124–6; and significance, 120; and the space of reasons, 126; and training, 146–7
 meiosis, 5, 221
 metaphor,; and assertion, 122–9; and attunement, 143–56; classification of theories of, 7–14; Cognitive account of, 187–91; commentaries on, 132–6; and conceptual art, 231; Conflated Sentences account of, 266–83; and context, 119–20, 249–50, 252–5, 263–6; dead, 182–202; and deceit, 242; Demonstrative account of, 249–66; depth and vividness of, 241–2; and elucidation, 135–7; and exemplars, 105–15; and the figurative, 5, 221–2, 231–9; as fundamental, 246–7; identification of, 10–11; Image account of, 26–8; and images, 12–13; Indirect Speech account of, 283–90; and individual words, 173–5; as a kind, 199–200, 221–31; as lumpy, 232–3; mixed, 178–82, 185–6; and the non-literal, 5; non-negotiable truths about, 7; and objects, 94–115;

- as ornamental, 246; and the Paraphrase requirement, 16–20, 129–36; philosophical accounts of, 3–5; and property sharing, 30–3, 37, 176–7, 264–6, 269–74; and qualification, 161–2; *see also* Semantic Descent account; and rationalization, 135–6; richness of, 239, 242–7; and seeing-as, 276–9; Semantic Descent account of (introduced), 94–105; and similarity, 30–3, 37, 176–7, 259–60, 264–6, 269–74; and synaesthesia, 234–9; and syntactic complexity, 157–78; and translation, 129–32; and the Transparency requirement, 21–35, 136–43; and the Truth requirement, 15–16, 121–9; units of, 159; visual, 227–31; and whole sentences, 173–5
- metaphor comprehension, speculation about, 141–3
- ‘metaphor for’, 173, 218–28, 261; and semantic descent, 219–28
- metaphor identification, 10, 117
- metaphor vs metaphorical, 159
- metaphorical complexity, 157–8
- metonymy, 5, 194–5, 221–6, 292
- mixed metaphor, 178–82, 185–6
- modularity and transparency, 23
- Monty Python, 185
- Moran, R., 13 n. 5, 17, 29 n. 22
- Mthat, 254–6, 258–61, 263–5, 289
- Nabokov, V., 72
- Nietzsche, F., 242–3
- No Message account, 9, 248, 287
- non-existence, 115
- Object account (introduced), 36–7; *see also* Semantic Descent account
- objects and context, 152–3
- objects as predicates, 59–72; *see also* qualification
- objects as words, 38–42, 62–85
- ontological asymmetry, 49–51
- oscillation (and bifurcation), 167–70
- parable, 5
- Paraphrase requirement, 16–20, 129–36, 155–6; and translation, 19, 129–32
- Peirce, C. S., 291–2
- pictures and truth, 229–30; and predicates, 65–6, 229–30
- Plato, 42, 77 n. 34
- Polshek, J., 69 n. 27
- polyseme, 140, 184
- pragmatics, 145–6; and semantics, 8–9
- Pred, 55–9
- predicates and subjects, 43–59
- predication; fair play for, 51–9; and exemplars, 105–15; a just-so story, 74–92, 200–2; and ontology, 50–1; and pictures, 65–6; Pred, 55–9; and qualification, 59–92; Strawson's model of, 53–4; task model of, 55–6
- primate (non-human) language, 89–90, 295–7
- properties; and concepts, 85–6, 88–9; and demonstratives, 102–3; and metaphor, 102–5; and semantic descent, 101–5
- property sharing and similarity, 269–74
- proto-predicates, 118–19, 121, 135
- proverbs, 221
- psycholinguistic speculations, 141–3
- qualification, 59–92; adicity of, 63; and attunement, 143–56, 178; as bringing to bear of information, 61; and categorization, 246–7, 78–82, 84; and communicative intentions, 71; and conceptual art, 231; and context, 63–4, 101, 119–20, 148–56, 249–50, 252–5, 263–6; and dead metaphor, 191–202; and determinate objects, 114–15; and fictional objects, 115; further examples of, 67–72; indeterminacy of, 149–56; initial examples of, 627; and interdependence with reference, 61–2, 294–7; and linguistic control, 153, 99–101; and literal predicates, 201; and ‘metaphor for’, 219; and natural meaning, 60–1, 71; and predication, 72–85, 143–6, 294

- ; and semantic descent, 94–115; and similarity, 30–3, 37, 176–7, 259–60, 264–6, 269–74; and syntactic complexity, 159–76; and visual metaphor, 227–31; vs metaphor, 161–2, 294
- Quine, W. V. O., 43, 52, 94, 116, 292
- quotation and down-arrow notation, 117, 261–3
- Ramachandran, V. S., 233–6
- rationalization of metaphor, 135–6, 138
- Recanati, F., 285 n. 18
- Ref, 51, 55
- reference, 55–60
- Rich, A. N., 236
- Roth, P., 158, 172 n. 10
- Russell, B., 262
- sad music, 238
- sample-series examples, 64–7
- saturated/unsaturated entities, 45–6
- Searle, J., 12, 14 n. 8, 19–20, 33, 121 n. 19, 183–5, 248, 267, 270 n. 9, 284–9
- seeing-as, 274–9
- Sellars, W., 126
- semantic ascent, 94, 116
- Semantic Descent account, 94–156; and adjectives, 173–5; and ambiguity, 140; and bifurcation, 166–73; and the Cognitive account, 188–91; and Conflated Sentences account, 279–83; and context, 148–56, 265–6; and dead metaphor, 191–202; and the Demonstrative account, 260–3, 265–6; and determinate objects, 114–15; down-arrow notation for, 116–19; and extended metaphor, 181–2; and fictional objects, 115; and the Image account, 138–43; and ‘like’, 207–17; and literal predicates, 201; and ‘metaphor for’, 219–28; and metonymy, 222–6; and mixed metaphor, 178–82; and objects, 94–115; and the Paraphrase requirement, 129–36; and properties, 101–5; and psychology, 112–13; and synecdoche, 222–6; and syntactic complexity, 159–76; and the Transparency requirement, 136–43; and the Truth requirement, 121–9; and visual metaphor, 227–31; without metaphor, 199; and whole sentences, 173–5
- ‘semantic’ in ‘semantic descent’, 117
- semantics and pragmatics, 8–9
- semantics vs pragmatics, 145–6
- Shakespeare, W., 8 n. 2, 161, 190, 240 n. 52, 267 n. 8, 268–9, 273, 278–9
- significance and context, 148–54
- significance and culture, 147–8
- similarity and metaphor, 30–3, 37, 102–5, 176–7, 259–60, 264–6, 269–74; and categorization, 80–1; and symmetry, 80–1
- Simile account, 202–18
- simile and symmetry, 211, 217
- Sperber, D., 284 n. 17
- Stern, J., 104 n. 7, 146, 150–1, 249–61, 262–6, 274, 289–90
- Strawson, P. F., 43–51, 53, 57, 86–90, 292
- subject and predicate, 42–59
- Suhamy, H., 221
- Swift, J., 52, 54, 74, 82–3
- symbols, 147, 292
- synaesthesia, 233–9; and metaphor, 234
- synecdoche, 5, 221–6, 292
- syntactic complexity, 157–78; and determinacy, 178
- systematicity of metaphor, 187
- Tarski, A., 59 n. 16
- Thomas, D., 158, 175
- thought and language, 295–7
- training and culture, 146–8
- training and meaning, 146–7
- translation (of metaphor), 129–32
- Transparency requirement, 21–35, 136–43, 285–6; and Alternative Message accounts, 33–4; and Content Sufficiency, 28–33; and the Image account, 26–8, 138–43; and literal meaning, 24–5, 137
- Travis, C., 145
- tropes, 5–6, 222–6
- Truth requirement, 15–16, 121–9
- Turner, M., 186
- Updike, J., 4, 245

visual metaphor, 227–31

White, R., 13 n. 4, 93, 157, 162, 164 n. 5, 165 n. 6, 166, 170
n. 9, 178, 181 n. 16, 182 n. 17, 189 n. 22, 240 n. 52,
266–9, 272–4, 276–83, 287

Wiggins, D., 42–51, 85–7

Wilson, D., 284 n. 17

Wittgenstein, L., 42, 52, 59 n. 16, 79 n. 35, 82, 123, 126 n.
21, 146, 238–9

Woolf, V., 218, 227–9, 292 n. 2

yellow vowels, 238–9

zeugma, 197