

The Philosophy of Philosophy

and
ton

Williamson

WILEY
Blackwell



The Philosophy of Philosophy

Second Edition

Timothy Williamson

WILEY Blackwell

This edition first published 2022
© 2022 John Wiley & Sons Ltd.

Edition History

Second Edition © 2022 John Wiley & Sons Ltd.
First Edition © 2007 Timothy Williamson, Published by Blackwell Publishing Ltd.

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, electronic, mechanical, photocopying, recording or otherwise, except as permitted by law. Advice on how to obtain permission to reuse material from this title is available at <http://www.wiley.com/go/permissions>.

The right of Timothy Williamson to be identified as the author of this work has been asserted in accordance with law.

Registered Offices

John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, USA
John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Office

111 River Street, Hoboken, NJ 07030, USA

For details of our global editorial offices, customer services, and more information about Wiley products visit us at www.wiley.com.

Wiley also publishes its books in a variety of electronic formats and by print-on-demand. Some content that appears in standard print versions of this book may not be available in other formats.

Limit of Liability/Disclaimer of Warranty

The contents of this work are intended to further general scientific research, understanding, and discussion only and are not intended and should not be relied upon as recommending or promoting scientific method, diagnosis, or treatment by physicians for any particular patient. In view of ongoing research, equipment modifications, changes in governmental regulations, and the constant flow of information relating to the use of medicines, equipment, and devices, the reader is urged to review and evaluate the information provided in the package insert or instructions for each medicine, equipment, or device for, among other things, any changes in the instructions or indication of usage and for added warnings and precautions. While the publisher and authors have used their best efforts in preparing this work, they make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Library of Congress Cataloging-in-Publication Data

A catalogue record for this book is available from the Library of Congress

Hardback ISBN: 9781119616672; ePDF ISBN: 9781119616696; epub ISBN: 9781119616726

Cover image: Portrait of Olga in an armchair Painting by Pablo Picasso (1881-1973) 1917,
JEAN-LOUIS JOSSE/Musee Picasso, Paris, France/Bridgeman Images
Cover design by Wiley

Set in 10.5/13 SabonLTStd by Integra Software Services Pvt. Ltd, Pondicherry, India

10 9 8 7 6 5 4 3 2 1

Contents

| | |
|---|---------------|
| <u><i>Preface to the Second Edition</i></u> | <i>xi</i> |
| <u><i>Preface to the First Edition</i></u> | <i>xxx</i> |
| <u><i>Acknowledgments</i></u> | <i>xxxiii</i> |
| | |
| Part I | 1 |
| <u>Introduction</u> | 3 |
| <u>1 The Linguistic Turn and the Conceptual Turn</u> | 12 |
| <u>2 Taking Philosophical Questions at Face Value</u> | 25 |
| <u>3 Metaphysical Conceptions of Analyticity</u> | 50 |
| <u>4 Epistemological Conceptions of Analyticity</u> | 75 |
| <u>5 Knowledge of Metaphysical Modality</u> | 136 |
| <u>6 Thought Experiments</u> | 181 |
| <u>7 Evidence in Philosophy</u> | 210 |
| <u>8 Knowledge Maximization</u> | 249 |
| <u>Afterword Must Do Better</u> | 280 |
| <u>Appendix 1 Modal Logic within Counterfactual Logic</u> | 295 |
| <u>Appendix 2 Counterfactual Donkeys</u> | 307 |
| | |
| Part II | 311 |
| <u>9 Widening the Picture</u> | 313 |
| <u>9.1 How Did We Get Here from There?</u> | |
| <u>The Transformation of Analytic Philosophy</u> | 313 |
| <u>9.2 Abductive Philosophy</u> | 351 |
| <u>9.3 Model-Building in Philosophy</u> | 372 |
| <u>9.4 Morally Loaded Cases in Philosophy</u> | 386 |

| | | |
|------|---|-----|
| 9.5 | Reply to Dennett and Kuznetsov on Abductive Philosophy | 401 |
| 9.6 | Reply to Kuznetsov and Stoljar on Model-Building in Philosophy | 404 |
| 10 | Experimental Philosophy | 406 |
| 10.1 | Reply to Weinberg | 406 |
| 10.2 | Philosophical Expertise and the Burden of Proof | 413 |
| 10.3 | On Joshua Alexander's <i>Experimental Philosophy: An Introduction</i> | 431 |
| 10.4 | Philosophical Criticisms of Experimental Philosophy | 440 |
| 10.5 | Reply to Dennett, Knobe, and Kuznetsov on "Philosophical Intuitions" | 464 |
| 11 | Naturalism | 467 |
| 11.1 | Reply to Kornblith | 467 |
| 11.2 | Reply to Stalnaker | 471 |
| 11.3 | Reply to Bianchi | 481 |
| 11.4 | What is Naturalism? | 484 |
| 11.5 | The Unclarity of Naturalism | 488 |
| 11.6 | On Penelope Maddy's <i>What Do Philosophers Do? Skepticism and the Practice of Philosophy</i> | 491 |
| 12 | Concepts, Understanding, Analyticity | 497 |
| 12.1 | Reply to Jackson | 497 |
| 12.2 | Reply to Boghossian | 502 |
| 12.3 | Reply to Peacocke | 512 |
| 12.4 | Reply to Mišćević | 520 |
| 12.5 | Reply to Smokrović | 529 |
| 12.6 | Reply to Trobok | 533 |
| 13 | Wittgensteinian Approaches | 538 |
| 13.1 | Reply to Moore | 538 |
| 13.2 | Reply to Horwich | 543 |
| 13.3 | Reply to Frascolla | 553 |
| 13.4 | Reply to Marconi | 556 |
| 13.5 | Reply to Tripodi | 560 |
| 13.6 | On Paul Horwich's <i>Wittgenstein's Metaphilosophy</i> | 563 |
| 14 | Miscellany | 569 |
| 14.1 | Reply to Ichikawa | 569 |
| 14.2 | Reply to Martin | 575 |

| | | |
|------|---|------------|
| 14.3 | <u>On Robert Brandom's <i>Reason in Philosophy: Animating Ideas</i></u> | <u>579</u> |
| 14.4 | <u>On Peter Unger's <i>Empty Ideas: A Critique of Analytic Philosophy</i></u> | <u>586</u> |
| 14.5 | Plato Goes Pop | 591 |
| 14.6 | Popular Philosophy and Populist Philosophy | 595 |
| | <u><i>Bibliography</i></u> | <u>598</u> |
| | <u><i>Index</i></u> | <u>619</u> |



Preface to the Second Edition

It was Marissa Koors, Philosophy editor at Wiley-Blackwell, who in 2018 proposed renewing *The Philosophy of Philosophy* in a second edition, with extra material on developments since 2007, when the book was first published. I liked the idea, without feeling tempted to rewrite the first edition. Since its publication, I have continued to stand behind all its main ideas and most of the details. In subsequent writings, I have further clarified and developed its lines of thought, responded to critics, and filled in omissions. However, those later pieces were scattered about, hard to survey and in some cases hard to find even for me, let alone anyone else. It may be helpful for readers to have all this material collected together into one volume, constituting a more comprehensive philosophy of philosophy, with replies to the sorts of questions and objections it tends to provoke.

My other projects delayed work on the second edition for over two years. This preface, written in the Oxford of 2020, under partial lockdown as a result of Covid-19, is an opportunity to look back, and forward, in briefly introducing the new material.

The most constructive additions are Sections 9.1–9.4, four essays that substantially extend the first edition's picture of philosophy, both its methods and its recent history. Each was written not so much as a contribution to an ongoing conversation as an attempt to start a new one. Those attempts already seem to be succeeding. Section 1, "Widening the picture," explains the topics of the new conversations, and how I came to be interested in them.

The other new sections, most of them quite short, and some of them quite polemical, were all written in something more like response-mode. Thus the distribution of topics in them is some evidence of what was happening in the philosophy of philosophy in the years after the publication of the first edition. The two response-mode sections

of full article length, Sections 10.2 and 10.4, are defenses of armchair philosophy against attacks from “experimental philosophers.” Of the shorter sections, twenty were my invited replies to book symposia on the first edition, in *Analysis*, *Philosophical Studies*, *Philosophy and Phenomenological Research*, *Analisi* (the bulletin of SIFA, the Italian Society of Analytic Philosophy), and the *Croatian Journal of Philosophy*, and to a symposium in *Epistemology and Philosophy of Science* (Moscow) on a paper in which I briefly summarized my updated view of philosophical methods (2019c).¹ Another five short sections originated as book reviews invited for *The Times Literary Supplement*, *Philosophy*, the *European Journal of Philosophy*, and *The Journal of Philosophy*. One commentary (14.5) originated in an invitation to review a large group of works of popular philosophy collectively for *The Times Literary Supplement*, another (14.6) in an invitation to contribute to the blog *Daily Nous*. Section 11.5 developed out of an invited reply for the *New York Times*’ philosophy blog “The Stone” to a defense of naturalism by Alex Rosenberg against my original post, out of which developed Section 11.4, itself provoked by “naturalist” responses to the first edition. I usually accept invitations to contribute to symposia on my books and articles, and to review books on topics on which I am currently working, though for many years my policy has been not write unsolicited replies to reviews or criticisms of my work; life is too short. Thus the balance of topics discussed in the additional response-mode sections is not an artefact of my selection.

All the sections have been written to be readable by themselves, which occasionally involves some local repetition. The response-mode material is overtly one-sided, since it includes only my half of each exchange – altogether, with nearly thirty philosophers, based in Australia, Canada, Croatia, Italy, Russia, the United Kingdom, and the United States. Of course, to judge properly whether I have been *fair* to my interlocutors, readers will have to read their side too. In any case, I am deeply grateful to all those who have spent so much time and effort carefully reading my work and articulating their responses.

In contrast to the first edition, the additional material is designed to be read selectively, according to the reader’s interests. It also varies in how wide a readership it was written for, depending on its original

¹ For citations in this form, “Williamson” is understood.

falsificationism of conjectures and refutations by counterexamples (thought experiments), though in a subtler and less direct way. One of our joint papers used an explicitly model-building methodology, and it was employed in an increasingly prominent role in some of my own publications from that period on.⁵ “Must Do Better,” the Afterword to the first edition, recommends the use of mathematical models to test philosophical ideas (293, this volume), though without discussing such methods in detail. Later reflection on the nature of progress in philosophy convinced me that, like progress in natural science, much of it takes the form of building better and better models of the phenomena under study, rather than discovering exceptionless universal laws, and that failure to recognize the model-building methodology is one of the reasons for widespread overestimation of the difference between philosophy and natural science. In that respect, the additional Section 9.3, “Model-Building in Philosophy,” goes far beyond the first edition, while Section 9.6 briefly considers a proposed alternative.⁶

A recent side interest, which played no role in the first edition, has been the surprisingly effective dialectical role of moral and political considerations in philosophical debates which seem to have nothing specifically to do with the moral or political – for example, over general relativism, general skepticism, and general internalism in epistemology. The story of how I first came to notice this phenomenon tempts me into a digression.

As a graduate student at Oxford, I used to attend meetings of the Radical Philosophy group, associated with the journal *Radical Philosophy*. In practice, what was philosophically radical about it was its rejection (and often ignorance) of analytic philosophy, in favor of just about anything which then counted as “continental” – they discussed Nietzsche, Saussure, Althusser, Derrida, the more arid parts of Foucault’s corpus, and so on, with varying degrees of reverence. The “analytic”-“continental” distinction cut at an obvious joint in the sociology of philosophy, however artificial it may have been in other respects. I experimented with those alternative traditions because I felt oppressed by the style and assumptions of the kind of analytic

⁵ For uses of model-building in my work see Shin and Williamson 1996, various passages in 2000a, and 2013c, 2014b, 2015a, 2019b, 2020b.

⁶ For a more “popular” account of model-building in philosophy see 2018a: 127–40.

philosophy then most fashionable in Oxford, and hoped that I might find different ideas for use in my own work. I didn't get much out of the experiments, though I enjoyed reading Nietzsche and Saussure. I came to realize that those who led the discussion often understood the obscure texts they talked about no more clearly than I did, although they certainly had a far more extensive acquaintance with them than mine, and were willing to "go on in the same way" as their authors. On the rare occasions when I asked a question or made an objection, they never seemed in danger of getting the point. There were one or two exceptions, fully open to rational discussion of ideas from both sides of the divide – one was Michael Rosen, now at Harvard. After I had left Oxford for my first proper teaching job, at Trinity College Dublin, I felt liberated to discover that what had really oppressed me about the then-predominant style of Oxford philosophy was not that it was too analytic but that it was not analytic *enough*. However, one of the things I did learn from my Oxford experience of Radical Philosophy was this: within such an intellectual world, much of the resistance to the relativist-sounding extremes of Post-Modernism came from Marxists and others on the far Left, who feared relativism as a threat to their political hopes. How far will those who view the case for revolution from a relativist stance commit to the revolutionary cause? In that world, objections to relativism from common sense, natural science, or logic had much less credibility. Later, while in Dublin (1980–1988), I was intrigued to hear from a talk by Richard Kearney (now at Boston College) of Richard Rorty describing absolutism about justice as much harder to give up than absolutism about truth. I was never tempted to give up either, but I could imagine how someone more concerned with morality and politics than with logic might feel that way.

I did nothing with those thoughts at the time, but they stayed with me. Much more recently, in responding to Paul Boghossian's epistemological internalism, I found myself objecting that it counts as justified (though false) a consistent neo-Nazi's belief that he ought to kill members of a target group, and wondering whether such a view would also count as justified (though wrong) his acting on that belief (Boghossian and Williamson 2020). That got me thinking more carefully about why emotive cases are dialectically effective, and whether invoking them is some kind of cheat. That is an obvious danger, especially in the current philosophical climate, where morally

or politically wrong-footing one's opponent is all too often used as a convenient excuse for not engaging properly with their arguments or objections. Nevertheless, I came to the conclusion that it *is* legitimate to use such examples in order to make vivid the practical consequences of a philosophical theory, especially one which had seemed to have none. The justification of belief and the justification of action should not be treated as orthogonal issues: the considerations for and against internalism are similar in the two cases, and after all the distinguishing mark of a belief is the agent's willingness to act on it. The additional Section 9.4, "Morally Loaded Cases in Philosophy," encapsulates my reflections on these issues.

The Preface to the first edition starts by expressing my long-held view that the self-images then salient for contemporary philosophy failed to fit its actual development over the preceding decades. The book aimed to help put that right. I had also long been aware of a related strangely growing gap in the historiography of analytic philosophy. When I started as an undergraduate at Oxford in 1973, historical narratives of analytic philosophy tended to stop the story around 1960. Naturally, I expected that, as time went on, the lag between the time of writing and the end of the period written about in narratives of analytic philosophy would remain roughly constant. It did not happen. Thirty years later, historical narratives of analytic philosophy *still* tended to stop the story around 1960. Although that generalization is not exceptionless, there really was very little serious historical work on post-1960 developments in analytic philosophy. The time lag was far longer than needed to gain some historical perspective on the past – it was far shorter for serious historical work on post-1960 (and indeed post-1989) developments in politics, society, and culture. Many younger philosophers felt that Saul Kripke, David Lewis, and others had effected a revolution in philosophy after 1960. I found it frustrating that no one seemed interested in achieving a proper historical understanding of so significant a change.

It was as though such a revolution was not *supposed* to happen. Whether historians of analytic philosophy preferred its logical positivist or its ordinary language strand, its predicted further development would not be in the direction of pre-Kantian metaphysics. From an older perspective, philosophers such as Kripke and Lewis looked like anomalies, anachronisms, to be swept away by the *zeitgeist*, unworthy of serious historical treatment. Instead, the opposite happened.

Their ways of doing philosophy gradually prevailed, to an extent increasingly hard to marginalize historically, whether one approved of them or not. The first edition of this book was obviously a product of that turn in philosophy, but did not say very much about its history.

Some years later, the historian of philosophy Miroslava Trajkovski encouraged me to give a talk at the University of Belgrade, to help bring later developments in analytic philosophy alive for her students by drawing on my personal acquaintance with many of the protagonists. I used the opportunity to reflect historically on the transition from linguistic philosophy to contemporary metaphysics, and describe how it felt to one person at the time. The result was my article “How did we get here from there? The transformation of analytic philosophy”, now included as the additional Section 9.1. It is not meant as a work of serious historical scholarship, but rather as a provocation to others to produce such works on post-1960 developments in analytic philosophy. Indeed, things had already begun to improve in that respect. Such historiography is now flourishing. For example, the massive influence of David Lewis has become well-recognized, and his key role in the history of post-1960 analytic philosophy is being analyzed in detail. After all, the period from 1960 to 2020 is just as long as that from 1900 to 1960, and just as deserving of historical study.^{7,8}

The reception of the first edition and of “How did we get here from there?” was in many ways gratifying. However, I will not resist one grumble. The experience brought home to me that not all historians of philosophy read a contemporary philosophical text with the professional accuracy or empathy one might expect. I give samples without naming names. Where I wrote “looked,” it was irritating to be read as if I had written “is”; I used the past (not present) tense and the verb “to look” (not “to be”) for a reason. It was irritating too to be read as if I must be using the word “analytic” in Kant’s sense, not in the clearly broader sense standard in analytic philosophy for the last half-century. It was also irritating when my deliber-

⁷ Incidentally, one of those now engaged in this much-needed work is Paolo Tripodi, with an earlier incarnation of whom I take issue in Chapter 13.5.

⁸ For a more general and “popular” discussion of the relation between philosophy and its history see 2018a: 98–110.

ately casual introduction of the phrase “armchair knowledge” for an overtly heterogeneous range of cases was read as aiming to replace the term “*a priori*” by a more precise substitute better suited to epistemological theorizing (171, this volume). Alas, no philosophical text is proof against determined attempts to interpret it to suit the interpreter’s purposes.

2. *Experimental philosophy*

The first edition treated another topic only briefly: the “negative program” of some “experimental philosophers” against “armchair philosophy.” It explained why their talk of “philosophical intuitions” failed to pick out a psychologically distinctive kind, and why thought experiments are not cognitively exceptional, as they assumed, but I did not engage with their texts in much detail.

However, the fashion for experimental philosophy was growing, and I often encountered (and still encounter) surprisingly crude misunderstandings of my objections to the negative program. Do the rejected obsolescent armchair methods include reading a philosophical text carefully and grasping its dialectical structure? In particular, many people took for granted that the book “defended philosophical intuitions,” when in fact it argued that thinking in terms of philosophical “intuitions” leads one hopelessly astray. I was also persistently classified as an “enemy of experimental philosophy,” despite having engaged in it myself (Bonini, Osherson, Viale, and Williamson 1999). Indeed, given the book’s keynote anti-exceptionalism about philosophy, it would have been absurd for me to argue that experimental results are *in principle* irrelevant to assessing the reliability of a philosophical method. But to assess it properly, you must first understand both what the method is and how it is being applied in particular cases. In practice, proponents of the negative program often – though not always – violated these conditions, either by seeing the methodological issues through the distorting lens of the category “philosophical intuitions,” or by making sundry naïve or impatient errors in handling the first-order philosophical issues themselves.

The negative program worried me because it had the potential to do serious damage to intellectual standards in philosophy – though its proponents’ intention was undoubtedly the opposite. Of course, no particular thought experiment is above criticism, just as no particular experiment in natural science is above criticism. But the negative

vicinity, sweep it off the table, and never return to it. That is naïve falsificationism at its worst.

A good strategy to deal with this problem is to hedge one's bets, by using more than one method. Each method acts as a potential corrective to the others. Where different methods converge on the same answer, our acceptance of it is correspondingly more robust. In particular, we can sometimes use both the case method and the method of model-building, neither having priority over the other. For example, I have used formal models of epistemic logic to argue that knowledge is not equivalent to justified true belief (as epistemologists traditionally use the word "justified"), the same conclusion normally reached in epistemology by Gettier-style thought experiments. Thus the two methods converge on the same answer.⁹ Although each method by itself may provide knowledge under normal conditions, in the long run we can expect more reliable results from using two or more methods to explore overlapping aspects of an issue and keep a check on each other.

One kind of normal human judgment about hypothetical cases which may sometimes go systematically wrong concerns *conditionals*. In *Suppose and Tell: The Semantics and Heuristics of Conditionals* (2020a), I explore what is arguably the primary human heuristic for cognitively assessing conditionals, a procedure which works well under most conditions, allowing us to extract and communicate valuable information stored in our dispositions to judgment about hypothetical cases, which are miniature thought experiments. This procedure is what we use to make judgments for or against the sample "if" sentences which provide most of the data for semantic and logical theories of conditionals in natural language. However, the heuristic *cannot* be fully reliable, for it is internally inconsistent. That explains why philosophers and linguists have had so much trouble agreeing on how "if" works. But the usual methods of experimental philosophy would not bring the limitations of the heuristic to light, if it is indeed a human universal, since all those who apply it are liable to the *same* errors. Rather, the heuristic's inconsistency is demonstrated by logical and

⁹ See 2013c, 2015a. These papers are not included in the present volume because they are mainly occupied with epistemological and technical issues, exemplifying rather than discussing philosophical methodology.

mathematical argument. The role of psychological experimentation lies elsewhere: in testing how far humans do indeed rely on that heuristic. That is a task for cognitive psychology, though not specifically for experimental philosophy.

In general, philosophy and cognitive psychology have much to learn from each other about the nature of human thought and its characteristic vices and virtues. Collaborations between philosophers and cognitive psychologists are likely to become increasingly fruitful, and trying to separate philosophy from psychology in the results may often be fruitless. Whether any of that should be described as “experimental philosophy” is another matter. Anti-exceptionalism about philosophy suggests that the psychology of human philosophical thinking is best understood as just a special case of the psychology of human thinking in general. Schematically: philosophers will have most to bring to their collaboration with psychologists by cultivating their distinctively philosophical skills, not by aping the psychologists, just as psychologists will have most to bring to the collaboration by cultivating their distinctively psychological skills, not by aping the philosophers – though, in a successful collaboration, the philosophers will surely learn lots of psychology and the psychologists lots of philosophy. One reason for the qualifier “schematically” is that there is already a continuum between “pure philosophy” and “pure psychology,” with different people at home on different points of the continuum. That is as it should be, and as it is on the continua between “pure philosophy” and “pure mathematics,” “pure physics,” “pure biology,” “pure computer science,” “pure linguistics,” “pure economics,” “pure history,” and so on. Philosophy has deep natural connections with many other disciplines; to give exclusive privileges to any one of them is to misunderstand the nature of philosophy.¹⁰

3. *Naturalism*

Of philosophers who self-identify as “naturalists,” the more extreme tend to dismiss *The Philosophy of Philosophy* as an anti-naturalist tract, while the more moderate tend to wonder why it does not make its implicit naturalism explicit. The first edition defends armchair

¹⁰ See 2018a: 111–26 for brief “popular” discussions of some close links between philosophy and various other disciplines.

philosophy against extreme naturalistic attacks, while also defending anti-exceptionalism about philosophy as much less different from other sciences in nature and methods than many philosophers like to think. It presents philosophy as an investigation of the same world which other sciences investigate too, and philosophical knowledge as the product of ordinary human cognitive capacities.

As best I can tell, there is an asymmetry between those who regard the book as implicitly naturalist and those who regard it as anti-naturalist: the former are more likely than the latter to have read it. After all, reading a book is an armchair method of learning what it says.

For the front cover of the first edition, I chose Picasso's "Portrait of Olga in an Armchair," because the sitter is a young woman, not the stereotypical philosopher in an armchair – an old man with a long beard and a pipe. The subliminal message was that armchair philosophy is not what you might think it is.

In a very loose sense of the term "naturalist," I probably count as one. The trouble is that the term is also often used much more narrowly, for one who takes the *natural* sciences (physics, chemistry, biology, ...) to provide the model which all other attempts at systematic inquiry should emulate in method. By that standard, even mathematics falls short, since it does not use observation or experiment in the intended sense, even though all the natural sciences rely on mathematics. It is the most obvious example of a science which is not a natural science in any distinctive sense. Another example, I suggest, is philosophy. The reliance on armchair methods is one of the most salient features of both mathematics and philosophy. That is not to deny the relevance of natural science to philosophy, or even to mathematics. It is just to insist that armchair methods have a central role to play in philosophy, and even more obviously in mathematics.

The second edition contains six short additional sections on naturalism, 11.1–11.6. Their main concerns are to separate extremist versions of naturalism from moderate ones, to emphasize the implausibility of the extremist versions, and to show that the moderate versions are fully compatible with armchair methods.

4. Concepts, understanding, analyticity

Some reactions to the book made me wish that I had been more explicit about my terminology. For example, I often used the words "concept" and "conceptual," but did little to define or clarify them.

The reason was that I borrowed those words from my *opponents*, primarily to articulate their views and arguments – to the effect that philosophy is in some distinctive sense a “conceptual” activity. I wanted to be fair to my opponents by not defining or clarifying the terms in ways which they might reject. Moreover, such views of philosophy come in numerous sub-varieties, which gloss the words in different ways, as I acknowledged (17, this volume). Since the same forms of argument often worked against different sub-varieties, I used the words “concept” and “conceptual” in a somewhat schematic way, to avoid unnecessary repetition.

The upshot of Chapter 4 is that there are no “conceptual” truths or connections in any sense helpful to my opponents. If one likes, one can define a “concept” to be the actual or potential meaning of a linguistic expression, which it shares with all synonymous expressions, appealing to whatever standard of sameness in meaning is made available by a well-developed semantic theory. However, I argued that such a standard will be too coarse-grained to serve my opponents’ purposes. For example, it will not make even the most elementary logical truths “conceptual” in any distinctive sense. Such conclusions should have made it clear that “concept” and “conceptual” were not load-bearing terms in my statements of my own positive views.

A little unwisely, I sometimes also wrote of applying “concepts” and of “conceptual” practices in stating my own views, not only in going along with my opponents’ ways of talking for the sake of argument. I could just as well have written instead of applying *words* and of *linguistic practices*. In those cases, the step of abstraction from linguistic expressions to concepts was idle. For instance, all the work can be done by the word “vixen” and the property of being a vixen, cutting out the useless middle man, the concept *vixen*.¹¹ In retrospect, I wish I had stuck to the more perspicuous metalinguistic formulations in stating my own views, and not muddied the waters by sentimentally continuing to employ the term “concept.”

For similar reasons, it would be more open to replace currently fashionable talk of “conceptual engineering” by talk of “linguistic engineering.” After all, our direct conscious and social control is of

¹¹ A more accurate statement replaces “the property of being a vixen” with a second-order analogue, as explained in 2013a.

linguistic practices rather than ways of thinking, and our indirect influence on the latter is typically through the former.¹²

For some readers, my use of the word “analytic” was also misleading, since they adhered to its older, historically and etymologically justified sense in which analytic truths are corollaries of conceptual analyses. On that view, “Vixens are female foxes” is both analytic and a conceptual truth, whereas “Red shades are not green” is not analytic but may still be a conceptual truth. I followed much current philosophical usage, which treats “analytic truth” and “conceptual truth” as interchangeable.

With these health warnings, I have left the terminology of the chapters from the first edition unchanged, since readers may wish to see how I originally put things, for purposes of comparison.

The six short additional sections, 12.1–12.6, are all replies to philosophers who took issue with the book on these topics.

5. *Other topics*

As a student at Oxford in the 1970s, my exposure to Wittgenstein’s influence helped me build up enough antibodies to resist it for a lifetime (see Section 9.1). Although his influence had greatly declined by the time I wrote the first edition (and has since declined further), he was still too salient a landmark to be ignored in a discussion of philosophical methodology, especially since in some respects my viewpoint stood directly opposite his. Responses to the first edition showed that his ideas were still widespread in the international community of philosophers. The six short additional sections, 13.1–13.6, all reply to philosophers whose approach to the philosophy of philosophy is strongly marked by Wittgenstein’s influence.

In the book, I did not intend to cast Wittgenstein, or anyone else, as the villain of the piece. Obviously, I am no Wittgenstein scholar; I am happy to leave detailed engagement with his texts to those with more interest in them. My primary interest has been in combating mistaken assumptions about philosophy widely held by living philosophers, without worrying too much about their historical origins. But philosophers with Wittgensteinian sympathies were strongly represented amongst the authors whom I was invited to respond to or review, perhaps because editors hoped for a lively debate.

¹² Both negative and positive associations of the word “control” are intended.

half of one of those exchanges, in a book symposium on the first edition. The later rounds are also relevant to the epistemological arguments of this book, although they are not specific to the epistemology of philosophy.

Chapter 7 of the first edition, on evidence in philosophy, in effect applies the general account of evidence defended in *Knowledge and its Limits* to the special case of philosophy. I have continued to uphold that account of evidence, though usually without special reference to philosophy (2021e).

I have also written for a much wider readership on what philosophy does and how (2018a, 2018d).¹⁴

The Preface to the first edition ends with an expression of my enjoyment in doing philosophy. I am happy to report that, fourteen years later, it continues to provide just as much pleasure.

¹⁴ The present text of the first edition corrects mistakes which escaped my proofreading, on pages 166, 180, and 306: thanks to Andrew Melnyk, Chi-Yen Liu, and David Etlin respectively. I have silently corrected a few similar mistakes in the added sections. I have also made various verbal adjustments in the added sections for the sake of smooth reading in the new context of the second edition.

Preface to the First Edition

This book grew out of a sense that contemporary philosophy lacks a self-image that does it justice. Of the self-images that philosophy inherited from the twentieth century, the most prominent – naturalism, the linguistic turn, post-modern irony, and so on – seemed obviously inadequate to most of the most interesting work in contemporary philosophy: as descriptions, false when bold, uninformative when cautious. Less prominent alternatives too seemed implausible or ill-developed. Although an adequate self-image is not a precondition of all virtue, it helps. If philosophy misconceives what it is doing, it is likely to do it worse. In any case, an adequate self-image is worth having for its own sake; we are not supposed to be leading the unexamined life. This is my attempt to do better.

I considered using the phrase “philosophical method” in the title, but decided against on the grounds that it seemed to promise something more like a recipe for doing philosophy than I believe possible. When asked for advice on some occasion, the Duke of Wellington is said to have replied “Sir, you are in a devilish awkward predicament, and must get out of it as best you can.” My advice would be scarcely more useful. At the crucial point, I can only say “Use your judgment.” The primary task of the philosophy of science is to understand science, not to give scientists advice. Likewise, the primary task of the philosophy of philosophy is to understand philosophy, not to give philosophers advice – although I have not rigorously abstained from the latter.

I also rejected the word “metaphilosophy.” The philosophy of philosophy is automatically part of philosophy, just as the philosophy of anything else is, whereas metaphilosophy sounds as though it might try to look down on philosophy from above, or beyond. One reason for the survival of implausible self-images of philosophy is

that they have been insufficiently scrutinized as pieces of philosophy. Passed down as though they were platitudes, they often embody epistemologically or logically naïve presuppositions. The philosophy of philosophy is no easier than the philosophy of science. And like the philosophy of science, it can only be done well by those with some respect for what they are studying.

The book makes no claim to comprehensiveness. For example, it does not engage in detail with critics of analytic philosophy who do not engage with it in detail. I preferred to follow a few lines of thought that I found more rewarding. I hope that philosophy as I have presented it seems worth doing and not impossibly difficult. At any rate, I enjoy it.



Acknowledgments

First come the acknowledgments from the first edition. My three Blackwell/Brown lectures, given at Brown University in September 2005, constituted the occasion for the book, although the material had evolved considerably since then. I thank both Blackwell Publishing and Brown University for the invitation and their generous hospitality. Jeff Dean at Blackwell was a helpful and supportive editor.

My further debts of gratitude are huge. An earlier version of some of the material was presented as the Jack Smart Lecture at the Australian National University in July 2005. Various later versions were presented as four Anders Wedberg Lectures at the University of Stockholm in April 2006, where the commentators were Kathrin Glüer-Pagin, Sören Häggqvist, Anna-Sara Malmgren, and Åsa Wikforss, as eight José Gaos Lectures at the Instituto de Investigaciones Filosóficas of the Universidad Nacional Autónoma de Mexico in September–October 2006, and as three Carl G. Hempel Lectures at Princeton University in December 2006. Other occasions on which the material in one form or another came under scrutiny included a week-long graduate course at the University of Bologna in May–June 2005, a week-long Kompaktseminar at the University of Heidelberg in February 2006, three lectures I gave as the Townsend Visitor in Philosophy at the University of California, Berkeley, in September 2006, a lecture and workshop at the University of Munich in June 2005, two lectures I gave as Tang Chun-I Visiting Professor at the Chinese University of Hong Kong in March 2007, and lectures at a graduate conference on epistemology at the University of Rochester in September 2004, where Richard Feldman was the commentator, the University of Arizona, Tucson, and the University of California, Los Angeles, and a meeting of the Aristotelian Society (my Presidential Address) in October 2004, a workshop on the epistemology of philosophy at

conference on philosophical knowledge. Chapters 7 and 8 derive from “Philosophical ‘Intuitions’ and Scepticism about Judgement,” *Dialectica* 58 (2004), pp. 109–53; the volume constitutes the proceedings of the workshop on intuition and epistemology at the University of Fribourg, Switzerland, in November 2002 (the talk I gave there is not recognizable in this book; I gave it to make myself think seriously about the topic). Chapter 7 in particular was greatly expanded; sections 1 and 7 were new; the probabilistic material in section 4 was expanded from pp. 683–5 of “Knowledge and Scepticism,” *The Oxford Handbook of Contemporary Philosophy*, edited by Frank Jackson and Michael Smith, (Oxford: Oxford University Press, 2005), pp. 681–700. The Afterword is a slightly modified version of “Must Do Better,” in *Truth and Realism*, edited by Patrick Greenough and Michael Lynch (Oxford: Oxford University Press: 2006), pp. 177–87; that volume constitutes the proceedings of the St. Andrews conference on meaning and truth.

The enlarged edition reprints the following previously published pieces, sometimes rearranged; for each, I was the sole author and all required permissions have been obtained, which I hereby acknowledge: “Replies to Kornblith, Jackson and Moore,” *Analysis Reviews*, 69 (2009): 125–35; “Replies to Ichikawa, Martin and Weinberg,” *Philosophical Studies*, 145 (2009): 465–76; “Plato Goes Pop,” *Times Literary Supplement*, 5529 (2009): 15; review of Robert Brandom, *Reason in Philosophy*, *Times Literary Supplement*, 5579 (2010): 22–3; “Philosophical Expertise and the Burden of Proof,” *Metaphilosophy*, 42 (2011): 215–29; “Reply to Peacocke,” “Reply to Boghossian,” “Reply to Stalnaker,” and “Reply to Horwich,” *Philosophy and Phenomenological Research*, 82 (2011): 481–7, 498–506, 515–23, and 534–42 respectively; “Three Wittgensteinians and a Naturalist on *The Philosophy of Philosophy*,” in Richard Davies (ed.), *Analisi: Annuario e Bollettino della Società Italiana di Filosofia Analitica (SIFA) 2011*, Milan: Mimesis (2011), 127–37; “What is Naturalism?” and “The Unclarity of Naturalism” in Matthew Haug (ed.), *Philosophical Methodology: The Armchair or the Laboratory?*, London: Routledge (2013): 29–31 and 36–8 respectively; review of Joshua Alexander, *Experimental Philosophy: An Introduction*, *Philosophy*, 88 (2013): 467–74; “Replies to Trobok, Smokrović, and Mišćević on the Philosophy of Philosophy,” *Croatian Journal of Philosophy*, 13 (2013): 49–64; review of Paul Horwich, *Wittgenstein’s Metaphilosophy*,

European Journal of Philosophy, 21 (2013): e7–e10; “How Did We Get Here From There? The Transformation of Analytic Philosophy,” *Belgrade Philosophical Annual*, 27 (2014): 7–37; review of Peter Unger, *Empty Ideas: A Critique of Analytic Philosophy*, *Times Literary Supplement*, 5833 (2015): 22–3; “Abductive Philosophy,” *Philosophical Forum*, 47 (2016): 263–80; “Philosophical Criticisms of Experimental Philosophy,” in Justin Sytsma and Wesley Buckwalter (eds.), *A Companion to Experimental Philosophy*, Oxford: Wiley Blackwell (2016): 22–36; “Model-building in Philosophy,” in Russell Blackford and Damien Broderick (eds.), *Philosophy’s Future: The Problem of Philosophical Progress*, Oxford: Wiley-Blackwell (2017): 159–73; review of Penelope Maddy, *What Do Philosophers Do? Skepticism and the Practice of Philosophy*, *Journal of Philosophy*, 114 (2017): 492–7; “Morally Loaded Cases in Philosophy,” *Proceedings and Addresses of the American Philosophical Association*, 93 (2019): 159–72; “Reply to Dennett, Knobe, Kuznetsov, and Stoljar on Philosophical Methodology,” *Epistemology and Philosophy of Science*, 56 (2019): 46–52; “Popular Philosophy and Populist Philosophy,” *Daily Nous* (2020), <http://dailynous.com/2020/06/08/popular-philosophy-populist-philosophy-guest-post-timothy-williamson>. Many of those pieces were indebted to various philosophers for their feedback. They are acknowledged at the relevant places in the text, since that is more informative.

I thank Marissa Koors, the acquisitions editor at Wiley-Blackwell, for first suggesting an expanded edition of the book to me, and for her patience, flexibility, and enthusiasm in facilitating its implementation, and Charlie Hamlyn, also at Wiley-Blackwell, for his detailed help in the later stages of the project.

Finally, as in the first edition, thanks above all to my wife Ana, who still does not let me forget what matters.



Part I

Introduction



constrain truth-directed inquiry. Linguistic or conceptual philosophers treat intuitions more sympathetically, as the deliverances of linguistic or conceptual competence. Of course, the appeal to intuitions also plays a crucial role in the overt methodology of other disciplines too, such as linguistics.

One main theme of this book is that the common assumption of philosophical exceptionalism is false. Even the distinction between the *a priori* and the *a posteriori* turns out to obscure underlying similarities. Although there are real methodological differences between philosophy and the other sciences, as actually practiced, they are less deep than is often supposed. In particular, so-called intuitions are simply judgments (or dispositions to judgment); neither their content nor the cognitive basis on which they are made need be distinctively philosophical. In general, the methodology of much past and present philosophy consists in just the unusually systematic and unrelenting application of ways of thinking required over a vast range of non-philosophical inquiry. The philosophical applications inherit a moderate degree of reliability from the more general cognitive patterns they instantiate. Although we cannot prove, from a starting-point a sufficiently radical skeptic would accept, that those ways of thinking are truth-conducive, the same holds of *all* ways of thinking, including the methods of natural science. That is the skeptic's problem, not ours. By more discriminating standards, the methodology of philosophy is not in principle problematic.

Some may wonder whether philosophy *has* a method to be studied, especially if it is as methodologically undistinctive as just suggested. Forget the idea of a single method, employed in all and only philosophical thinking. Still, philosophers use methods of various kinds: they philosophize in various ways. A philosophical community's methodology is its repertoire of such methods. The word "method" here carries no implication of a mechanically applicable algorithm, guaranteed to yield a result within a finite time. On this loose understanding of what a methodology is, it is disingenuous for a philosopher to claim to have none.

Another main theme of this book is that the differences in subject matter between philosophy and the other sciences are also less deep than is often supposed. In particular, few philosophical questions are conceptual questions in any distinctive sense, except when philosophers choose to ask questions about concepts, as they may but need

not do. Philosophical questions are those philosophers are disposed to ask, which in turn tend, unsurprisingly, to be those more amenable to philosophical than to other ways of thinking; since the philosophical ways of thinking are not different in kind from the other ways, it is equally unsurprising that philosophical questions are not different in kind from other questions. Of course, philosophers are especially fond of abstract, general, necessary truths, but that is only an extreme case of a set of intellectual drives present to some degree in all disciplines.

In most particular cases, philosophers experience little difficulty in recognizing the difference between philosophy and non-philosophy. Being philosophers, they care about the difference, and have a professional temptation to represent it as a deep philosophical one. But just about every institutionally distinct discipline acquires a professional identity, and its practitioners experience little difficulty in recognizing the difference between what “we” do and what “they” do in most particular cases. They care about the difference, and have a professional temptation to represent it in the terms of their own discipline. But such temptations can be resisted. The distinction between the Department of Philosophy and the Department of Linguistics or the Department of Biology is clearer than the distinction between philosophy and linguistics or biology; the philosophy of language overlaps the semantics of natural languages and the philosophy of biology overlaps evolutionary theory.

The unexceptional nature of philosophy is easier to discern if we avoid the philistine emphasis on a few natural sciences, often imagined in crudely stereotyped ways that marginalize the role of armchair methods in those sciences. Not all science is natural science. Whatever crude empiricists may say, mathematics is a science if anything is; it is done in an armchair if anything is. In no useful sense are mathematical questions conceptual questions. If mathematics is an armchair science, why not philosophy too?

Most philosophers are neither crude rationalists nor crude empiricists nor, these days, linguistic or conceptual philosophers. Many would accept the theses just enunciated about the methodology and subject matter of philosophy. But a third theme of this book is that the current philosophical mainstream has failed to articulate an adequate philosophical methodology, in part because it has fallen into

the classic epistemological error of psychologizing the data. For example, our evidence is sometimes presented as consisting of our intuitions: not their content, since it is allowed that some of our intuitions may be false, but rather our psychological states of having those intuitions. We are then supposed to infer to the philosophical theory that best explains the evidence. But since it is allowed that philosophical questions are typically not psychological questions, the link between the philosophical theory of a non-psychological subject matter and the psychological evidence that it is supposed to explain becomes problematic: the description of the methodology makes the methodology hard to sustain. Again, philosophy is often presented as systematizing and stabilizing our beliefs, bringing them into reflective equilibrium: the picture is that in doing philosophy what we have to go on is what our beliefs currently are, as though our epistemic access were only to those belief states and not to the states of the world that they are about. The picture is wrong; we frequently have better epistemic access to our immediate physical environment than to our own psychology. A popular remark is that we have no choice but to start from where we are, with our current beliefs. But where we are is not only having various beliefs about the world; it is also having significant knowledge of the world. Starting from where we are involves starting from what we already know, and the goal is to know more (of course, how much more we come to know cannot be measured just by the number of propositions learnt). To characterize our method as one of achieving reflective equilibrium is to fail to engage with epistemologically crucial features of our situation. Our understanding of philosophical methodology must be rid of internalist preconceptions.

Philosophical errors distort our conception of philosophy in other ways too. Confused and obscure ideas of conceptual truth create the illusion of a special domain for philosophical investigation. Similarly, although perception clearly involves causal interaction between perceiver and perceived, crudely causal accounts of perceptual knowledge that occlude the contribution of background theory create the illusion of a contrast between world-dependent empirical beliefs and world-independent philosophical theory.

Clearly, the investigation of philosophical methodology cannot and should not be philosophically neutral. It is just more philosophy,

turned on philosophy itself. We have the philosophy of mathematics, the philosophy of physics, the philosophy of biology, the philosophy of economics, the philosophy of history; we also need the philosophy of philosophy.

The rethinking of philosophical methodology in this book involves understanding, at an appropriate level of abstraction, how philosophy is actually done. Philosophers of science know the dangers of moralizing from first principles on how a discipline should ideally be pursued without respecting how it currently is pursued; the same lesson applies to the philosophy of philosophy. The present opposition to philosophical exceptionalism is far from involving the idea that philosophers should model themselves on physicists or biologists. The denial that philosophical questions are conceptual questions is quite compatible with a heavy emphasis on issues of semantic structure in philosophical discussion, for the validity or otherwise of philosophical reasoning is often highly sensitive to delicate aspects of the semantic structure of premises and conclusion: to make our reasoning instruments more reliable, we must investigate those instruments themselves, even when they are not the ultimate objects of our concern.

That philosophy *can* be done in an armchair does not entail that it *must* be done in an armchair.¹ This book raises no objection to the idea that the results of scientific experiments are sometimes directly relevant to philosophical questions: for example, concerning the philosophy of time. But it is a fallacy to infer that philosophy can nowhere usefully proceed until the experiments are done. In this respect, philosophy is similar to mathematics. Scientific experiments can be relevant to mathematical questions. For instance, a physical theory may entail that there are physically instantiated counterexamples to a mathematical theory. A toy example: one can specify in physical terms what it takes to be an inscription (intended or unintended) in a given font of a proof of “ $0 = 1$ ” in a given formal system of Peano Arithmetic; a physical theory could predict that an event of a specified physically possible type would cause there to be

¹ In this respect Hilary Kornblith seems to misunderstand the claim that philosophy can be done in an armchair (2006: 19). I have even dabbled in experimental philosophy myself (Bonini, Osherson, Viale and Williamson 1999).

such an inscription. Less directly, psychological experiments might in principle reveal levels of human unreliability in proof-checking that would undermine current mathematical practice. To conclude on that basis alone that mathematics should become an experimental discipline would be hopelessly naïve. In practice, most of mathematics will and should remain an armchair discipline, even though it is not in principle insulated from experimental findings, because armchair methods, specifically proof, remain by far the most reliable and efficient available. Although the matter is less clear-cut, something similar may well apply to many areas of philosophy, for instance, philosophical logic. In particular, on the account in this book, the method of conducting opinion polls among non-philosophers is not very much more likely to be the best way of answering philosophical questions than the method of conducting opinion polls among non-physicists is to be the best way of answering physical questions.

Although this book is a defense of armchair philosophy, it is not written in a purely conservative spirit. Our ideas about philosophical methodology, however inchoate, are liable to influence the methodology we actually employ; bad ideas about it are liable to tilt it in bad directions. A reasonable hypothesis is that our current methodology is good enough to generate progress in philosophy, but not by much: ten steps forward, nine steps back. Nevertheless, we can improve our performance even without radically new methods. We need to apply the methods we already have with more patience and better judgment. A small increase in accuracy of measurement may enable scientists to tackle problems previously beyond reach, because their data lacked sufficient resolution. Similarly, small improvements in accepted standards of reasoning may enable the philosophical community to reach knowledgeable agreement on the status of many more arguments. Such incremental progress in philosophical methodology is a realistic prospect, for current standards in the profession exhibit large variations significantly correlated with differences between graduate schools. Philosophical methodology can be taught – mainly by example, but fine-tuning by explicit precept and discussion also makes a difference. For instance, the level of rigor in philosophical statement and argument which Frege achieved only by genius (with a little help from his mathematical training) is now available to hundreds of graduate students every year: and we know how to do even better. That is not to imply, of course, that we must strive for maximum

1

The Linguistic Turn and the Conceptual Turn

The Linguistic Turn is the title of an influential anthology edited by Richard Rorty, published in 1967. He credited the phrase to Gustav Bergmann (Bergmann 1964: 3; Rorty 1967: 9). In his introduction, Rorty (1967: 3) explained:

The purpose of the present volume is to provide materials for reflection on the most recent philosophical revolution, that of linguistic philosophy. I shall mean by “linguistic philosophy” the view that philosophical problems are problems which may be solved (or dissolved) either by reforming language, or by understanding more about the language we presently use.

“The linguistic turn” has subsequently become the standard vague phrase for a diffuse event – some regard it as *the* event – in twentieth-century philosophy, one not confined to signed-up linguistic philosophers in Rorty’s sense. For those who took the turn, language was somehow the central theme of philosophy.

The word “theme” is used with deliberate vagueness. It does not mean “subject matter,” for the linguistic turn was not the attempted reduction of philosophy to linguistics. The theme of a piece of music is not its subject matter. Those who viewed philosophy as an activity of dispelling confusions of linguistic origin did not see it as having a subject matter in the sense in which a science has a subject matter. But merely to regard linguistic analysis as one philosophical method among many is not yet to have taken the linguistic turn, for it is not yet to regard language as central. We will be more precise below.

There is an increasingly widespread sense that the linguistic turn is past. We will ask how far the turn has been, or should be, reversed.

Language has been regarded as central to philosophy in many different ways, which cannot all be treated together. A history of the many different forms that the linguistic turn took would be a history of much of twentieth-century philosophy. That is a task for another book, by another author. Self-indulgently, I will use a thin slice through history to introduce the contemporary issues by briefly considering some of my predecessors in the Wykeham Chair of Logic at Oxford.

A. J. Ayer was the first holder of the Chair to take the linguistic turn.¹ In 1936, back from Vienna and its Circle but not yet in the Chair, he announced an uncompromisingly formal version of linguistic philosophy:

[T]he philosopher, as an analyst, is not directly concerned with the physical properties of things. He is concerned only with the way in which we speak about them. In other words, the propositions of philosophy are not factual, but linguistic in character – that is, they do not describe the behaviour of physical, or even mental, objects; they express definitions, or the formal consequences of definitions. (Ayer 1936: 61–2)

Ayer traced his views back ultimately to the empiricism of Berkeley and Hume (Ayer 1936: 11). His contrast between definitions of words and descriptions of objects is, roughly, the linguistic analogue of Hume's contrast between relations of ideas and matters of fact. For an empiricist, the *a priori* methods of philosophy cannot provide us with knowledge of synthetic truths about matters of fact ("the behaviour of physical, or even mental, objects"); they yield only analytic truths concerning relations of ideas ("definitions, or the formal consequences of definitions"). A rather traditional empiricism later overshadowed the linguistic theme in Ayer's work.

Ayer was the predecessor of Sir Michael Dummett in the Wykeham Chair. Dummett gave a much-cited articulation of the linguistic turn, attributing it to Frege:

Only with Frege was the proper object of philosophy finally established: namely, first, that the goal of philosophy is the analysis of the

¹ Ayer's three immediate predecessors were John Cook Wilson, H. H. Joachim and H. H. Price.

structure of *thought*; secondly, that the study of *thought* is to be sharply distinguished from the study of the psychological process of *thinking*; and, finally, that the only proper method for analysing thought consists in the analysis of *language*. . . . [T]he acceptance of these three tenets is common to the entire analytical school. (Dummett 1978: 458)

On this view, thought is essentially expressible (whether or not actually expressed) in a public language, which filters out the subjective noise, the merely psychological aspects of thinking, from the inter-subjective message, that which one thinks. Dummett's own corpus constitutes one of the most imposing monuments of analytic philosophy as so defined. Unlike Ayer, he does not describe philosophical claims as definitions. Unlike Rorty, he characterizes the linguistic turn as involving distinctive claims about the subject matter of philosophy, not only about its method. On Dummett's view, Frege's insight replaced epistemology by philosophy of language as first philosophy. But this methodological innovation is supposed to be grounded in the account of the proper object of philosophy.

Elsewhere, Dummett makes clear that he takes this concern with language to be what distinguishes "analytical philosophy" from other schools (1993: 4). His account of its inception varies slightly. At one points (1993: 5), he says: "[A]nalytical philosophy was born when the 'linguistic turn' was taken. This was not, of course, taken uniformly by any group of philosophers at any one time: but the first clear example known to me occurs in Frege's *Die Grundlagen der Arithmetik* of 1884." Later (1993: 27), we read: "If we identify the linguistic turn as the starting-point of analytical philosophy proper, there can be no doubt that, to however great an extent Frege, Moore and Russell prepared the ground, the crucial step was taken by Wittgenstein in the *Tractatus Logico-Philosophicus* of 1922." Presumably, in Frege the linguistic turn was a fitful insight, in Wittgenstein, a systematic conception.

That "analytical philosophers" in Dummett's sense coincide with those usually classified as such is not obvious. Some kind of linguistic turn occurred in much of what is usually called "continental [supposedly non-analytic] philosophy." That Jacques Derrida did not subscribe in his own way to Dummett's three tenets is unclear: if some stretching of terms is required, it is for the later Wittgenstein

too. Conversely, Bertrand Russell did not subscribe to the three tenets, although often cited as a paradigm “analytical philosopher.” Over the past 20 years, fewer and fewer of those who would accept the label “analytic philosophy” for their work would also claim to take the linguistic turn (I am not one of those few). Even philosophers strongly influenced by Dummett, such as Gareth Evans, Christopher Peacocke, and John Campbell, no longer give language the central role he describes. For Dummett, they belong to a tradition that has grown out of “analytical philosophy” without themselves being “analytical philosophers” (1993: 4–5). In effect, they aimed to analyze thought directly, without taking a diversion through the analysis of language. In the 1980s it became commonplace in some circles to suggest that the philosophy of mind had displaced the philosophy of language in the driving seat of philosophy.

For philosophers of mind who accepted Jerry Fodor’s (1975) influential hypothesis of a language of thought, the priority of thought to public language did not imply the priority of thought to all language, since thought itself was in a language, the brain’s computational code. In principle, someone might combine that view with Dummett’s three tenets of analytic philosophy, contrary to Dummett’s intention; he did not mean a private language. Moreover, the first-personal inaccessibility of the language of thought makes such a version of the linguistic turn methodologically very different from the traditional ones.

For those who deny the methodological priority of language to thought, the minimal fallback from Dummett’s three tenets is to reject the third but maintain the first two. They assert that the goal of philosophy is the analysis of the structure of thought, and that the study of thought is to be sharply distinguished from the study of the psychological process of thinking, but deny that the only proper method for analysing thought consists in the analysis of language. If thought has constituents, we may call them “concepts.” On this view, concepts take the place of words in Dummett’s analytical philosophy.

In practice, linguistic philosophers were often happy enough to speak of concepts rather than words, for they regarded a concept as what synonymous expressions had in common; their primary interest was in the features common to synonyms, not in the differences between them. It is therefore not too misleading to describe as *conceptual philosophers* those who accept Dummett’s first two tenets –

that the goal of philosophy is the analysis of the structure of thought, and that the study of thought is to be sharply distinguished from the study of the psychological process of thinking – whether or not they accept the third. We may also describe them as doing *conceptual philosophy*, and as having taken the *conceptual turn*.

The conceptual turn constitutes a much broader movement than the linguistic turn. It is neutral over the relative priority of language and thought. We think and talk about things – truly or falsely depending on whether they are or are not as we think or say they are. The aboutness of thought and talk is their *intentionality*; the conceptual turn puts intentionality at the centre of philosophy. This terminology indicates how little the conceptual turn is confined to what would ordinarily be called “analytic philosophy.” The phenomenological tradition may constitute another form of the conceptual turn. In the hermeneutic study of interpretation and various shades of postmodernist discourse about discourse the conceptual turn takes a more specifically linguistic form.

Have we stretched our terms so far that all philosophy is conceptual philosophy? No. On a natural view, concepts constitute only a small fraction of a largely mind-independent reality. That the goal of philosophy is in some sense to analyze that small fraction is no platitude. To put it very schematically, let *absolute idealism about the subject matter of philosophy* be the view that philosophy studies only concepts, in contrast to *ontological absolute idealism*, the wilder view that only concepts exist.² Although absolute idealism about the subject matter of philosophy does not entail ontological absolute idealism, why should we accept absolute idealism about the subject matter of philosophy if we reject ontological absolute idealism? Of course, we might reject absolute idealism about the subject matter of philosophy while nevertheless holding that the correct method for philosophy is to study its not purely conceptual subject matter by studying concepts of that subject matter. This methodological claim will be considered later; for present purposes, we merely note how much weaker it is than those formulated by Ayer and Dummett.

The claim that concepts constitute only a small fraction of reality might be opposed on various grounds. Recall that concepts were

² The “absolute” is to distinguish these forms of idealism from the corresponding “subjective” forms, in which concepts are replaced by psychological processes.

principle of being individually thought of? Although we can think of them collectively – for example, as elusive objects – that is not to single out any one of them in thought. Can we be sure that ordinary material objects do not consist of clouds of elusive sub-sub-atomic particles? We might know them by their collective effects while unable to think of any single one of them. The general question whether there can be elusive objects looks like a good candidate for philosophical consideration. Of course, McDowell does not intend the conceptual to be limited by the merely medical limitations of human beings, but the elusiveness may run deeper than that: the nature of the objects may preclude the kind of separable causal interaction with complex beings that isolating them in thought would require. In Fregean terminology again, a sense is a mode of presentation of a referent; a mode of presentation of something is a way of presenting it to a possible thinker, if not an actual one; for all McDowell has shown, there may be necessary limitations on thinking.⁴ Although elusive objects belong to the same very general ontological category of objects as those we can single out, their possibility still undermines McDowell's claim that we cannot make "interesting sense" of the idea of something outside the conceptual realm (1994: 105–6). We do not know whether there actually are elusive objects. What would motivate the claim that there are none, if not some form of idealism very far from McDowell's intentions? We should adopt no conception of philosophy that on methodological grounds excludes elusive objects.⁵

Suppose, just for the sake of argument, that there are no elusive objects. That by itself would still not vindicate a restriction of philosophy to the conceptual, the realm of sense or thought. The practitioners of any discipline have thoughts and communicate them,

⁴ McDowell's invocation of humility (1994: 40) addresses contingent limitations, not necessary ones.

⁵ Mark Johnston (1993: 96–7) discusses "the Enigmas, entities essentially undetectable by us." He stipulates that they are collectively as well as individually undetectable; thus our elusive objects need not be his Enigmas. If we cannot have good evidence that there are no Enigmas, it may well be a waste of time to worry whether there are Enigmas. But it would not follow that it is a waste of time to worry whether there *can be* Enigmas. Their definition does not rule out knowledge of the possibility of such things; such knowledge may itself be philosophically useful (indeed, Johnston uses it for his philosophical purposes).

but they are rarely studying those very thoughts: rather, they are studying what their thoughts are about. Most thoughts are not about thoughts. To make philosophy the study of thought is to insist that philosophers' thoughts should be about thoughts. It is not obvious why philosophers should accept that restriction.

Even within what is usually considered analytic philosophy of mind, much work violates the two tenets of conceptual philosophy. Naturalists hold that everything is part of the natural world, and should be studied as such; many of them study thought as part of the natural world by not sharply distinguishing it from the psychological process of thinking. Those who study sensations or qualia without treating them as intentional phenomena are not usually attempting to analyze the structure of thought; their interest is primarily in the nature of the sensations or qualia themselves, not in our concepts of them. Even when the question of veridicality arises, it is not always conceded that there are structured thoughts: some philosophers claim that perception has a conceptually unstructured content that represents the environment as being a certain way. Their interest is in the nature of the nonconceptual content itself, not just in our concept of it.

Despite early hopes or fears, philosophy of mind has not come to play the organizing role in philosophy that philosophy of language once did. No single branch of philosophy does: philosophy is no more immune than other disciplines to increasing specialization. Nor is any one philosophical method currently treated as a panacea for philosophical ills, with consequent privileges for its home branch. Once we consider other branches of philosophy, we notice much more philosophizing whose primary subject matter is not conceptual.

Biology and physics are not studies of thought. In their most theoretical reaches, they merge into the philosophy of biology and the philosophy of physics. Why then should philosophers of biology and philosophers of physics study only thought? Although they sometimes study what biologists' and physicists' concepts are or should be, sometimes they study what those concepts are concepts of, in an abstract and general manner. If the conceptual turn is incompatible with regarding such activities as legitimately philosophical, why take the conceptual turn?

There is a more central example. Much contemporary metaphysics is not primarily concerned with thought or language at all. Its goal

is to discover what fundamental kinds of things there are and what properties and relations they have, not to study the structure of our thought about them – perhaps we have no thought about them until it is initiated by metaphysicians. Contemporary metaphysics studies substances and essences, universals and particulars, space and time, possibility and necessity. Although nominalist or conceptualist reductions of all these matters have been attempted, such theories have no methodological priority and generally turn out to do scant justice to what they attempt to reduce.

The usual stories about the history of twentieth-century philosophy fail to fit much of the liveliest, exactest, and most creative achievements of the final third of that century: the revival of metaphysical theorizing, realist in spirit, often speculative, sometimes commonsensical, associated with Saul Kripke, David Lewis, Kit Fine, Peter van Inwagen, David Armstrong and many others: work that has, to cite just one example, made it anachronistic to dismiss essentialism as anachronistic.⁶ On the traditional grand narrative schemes in the history of philosophy, this activity must be a throwback to pre-Kantian metaphysics: it ought not to be happening – but it is. Many of those who practice it happily acknowledge its continuity with traditional metaphysics; appeals to the authority of Kant, or Wittgenstein, or history, ring hollow, for they are unbacked by any argument that has withstood the test of recent time.

One might try to see in contemporary metaphysics a Quinean breakdown of divisions between philosophy and the natural sciences. But if it is metaphysics naturalized, then so is the metaphysics of Aristotle, Descartes, and Leibniz. Armchair argument retains a central role, as do the modal notions of metaphysical possibility and necessity. Although empirical knowledge constrains the attribution of essential properties, results are more often reached through a subtle interplay of logic and the imagination. The crucial experiments are thought experiments.

Might the contrast between the new-old metaphysics and the conceptual turn be less stark than it appears to be? Contemporary metaphysicians firmly resist attempts to reconstrue their enterprise as

⁶ On essentialism see, for example, Kripke (1980), French, Uehling, and Wettstein (1986), Fine (1994, 1995) and Wiggins (2001). For a good statement of the outlook of contemporary metaphysics see Zimmerman (2004).

the analysis of thought – unlike Sir Peter Strawson, who defined his “descriptive metaphysics” as “content to describe the actual structure of our thought about the world” (1959: 9). But can one reflect on concepts without reflecting on reality itself? For the aboutness of thought and talk is their very point. This idea has been emphasized by David Wiggins, Dummett’s successor and my predecessor in the Wykeham Chair, and author of some of the most distinguished essentialist metaphysics, in which considerations of logic and biology harmoniously combine. Wiggins (2001: 12) writes: “Let us forget once and for all the very idea of some knowledge of language or meaning that is not knowledge of the world itself.”

Wiggins is not just stating the obvious, that language and meaning are part of the world because everything is part of the world. Rather, his point is that in defining words – natural kind terms, for instance – we must point at real specimens. What there is determines what there is for us to mean. In knowing what we mean, we know something about what there is. That prompts the question how far the analysis of thought or language can be pursued autonomously with any kind of methodological priority.

Dummett claimed not that the traditional questions of metaphysics cannot be answered but that the way to answer them is by the analysis of thought and language. For example, in order to determine whether there are numbers, one must determine whether number words such as “7” function semantically like proper names in the context of sentences uttered in mathematical discourse. But what is it so to function? Although devil words such as “Satan” appear to function semantically like proper names in the context of sentences uttered in devil-worshipping discourse, one should not jump to the conclusion that there are devils. However enthusiastically devil-worshippers use “Satan” as though it referred to something, that does not make it refer to something. Although empty names *appear* to function semantically like referring names in the context of sentences uttered by those who believe the names to refer, the appearances are deceptive. “Satan” refers to something if and only if some sentence with “Satan” in subject position (such as “Satan is self-identical”) expresses a truth, but the analysis of thought and language is not the best way to discover whether any such sentence does indeed express a truth. Of course, what goes for “Satan” may not go for “7.” According to some neo-logicians, “7 exists” is an analytic truth (what

Ayer might have called a formal consequence of definitions), which “Satan exists” does not even purport to be. Such a claim needs the backing of an appropriate theory of analyticity.

After this preliminary sketch, it is time to get down to detailed work. The next three chapters examine different forms of the linguistic or conceptual turn. Chapter 2 uses a case study to consider in a microcosm the idea that philosophers’ questions are implicitly about language or thought when they are not explicitly so. Chapters 3 and 4 assess a wide range of versions of the idea that the armchair methodology of philosophy is grounded in the analytic or conceptual status of a core of philosophical truths, which need not be *about* language or thought, even implicitly. In each case the upshot is negative. Although philosophers have more reason than physicists to consider matters of language or thought, philosophy is in no deep sense a linguistic or conceptual inquiry, any more than physics is. But it does not follow that experiment is an appropriate primary method for philosophy. Similar arguments suggest that mathematics is in no deep sense a linguistic or conceptual inquiry, yet experiment is not an appropriate primary method for mathematics. The second half of the book develops an alternative conception of philosophy, on which a largely armchair methodology remains defensible, as it does for mathematics.

From this perspective and that of many contemporary philosophers, the conceptual turn and *a fortiori* the linguistic turn look like wrong turnings. It is pointless to deny that such philosophers are “analytic,” for that term is customarily applied to a broad, loose tradition held together by an intricate network of causal ties of influence and communication, not by shared essential properties of doctrine or method: what do Frege, Russell, Moore, Wittgenstein, Carnap, Ayer, Quine, Austin, Strawson, Davidson, Rawls, Williams, Anscombe, Geach, Armstrong, Smart, Fodor, Dummett, Wiggins, Marcus, Hintikka, Kaplan, Lewis, Kripke, Fine, van Inwagen and Stalnaker all have in common to distinguish them from all the non-analytic philosophers? Many who regard the linguistic and conceptual turns as serious mistakes have ties of influence and communication that put them squarely within that tradition. “Analytic philosophy” is a phrase in a living language; the attempt to stipulate a sense for it that excludes many of the philosophers just listed will achieve nothing but brief terminological confusion.

it was dry; other measures would have been equally inconclusive. We have no idea of any investigative procedure that would have resolved the issue. It was a borderline case. No urgent practical purpose compels us to ask whether Mars was dry then, but only a limited proportion of thought and talk in any human society is driven by urgent practical purposes. We should like to know the history of Mars. When necessary, we can always use words other than “dry.” Nevertheless, we reflect on the difficulty of classifying Mars as dry or as not dry at those intermediate times, even given exact measurements. We may wonder whether it was either. We ask ourselves:

Was Mars always either dry or not dry?

Henceforth I will refer to that as *the original question*. More precisely, I will use that phrase to designate that interrogative sentence, as used in that context (the word “question” can also be applied to what interrogative sentences express rather than the sentences themselves).

The original question is at least proto-philosophical in character. It is prompted by a difficulty both hard to identify and hard to avoid that we encounter in applying the distinctions in our repertoire. It hints at a serious threat to the validity of our most fundamental forms of deductive reasoning. Philosophers disagree about its answer, on philosophical grounds explored below. A philosophical account of vagueness that does not tell us how to answer the original question is thereby incomplete. Without an agreed definition of “philosophy,” we can hardly expect to *prove* that the original question or any other is a philosophical question; but when we discuss its answer, we find ourselves invoking recognizably philosophical considerations. More simply, I’m a philosopher, I find the original question interesting, although I think I know the answer, and I have no idea where one should go for an answer to it, if not to philosophy (which includes logic). But before we worry about the answer, let us examine the original question itself.

The question queries just the supposition that Mars was always either dry or not dry, which we can formalize as a theorem of classical logic, $\forall t (\text{Dry}(m, t) \vee \neg \text{Dry}(m, t))$.² In words: for every time t ,

² Classical logic is the standard logic of expressions such as “every,” “either . . . or . . .” and “not” on the assumption that there is a mutually exclusive, jointly exhaustive dichotomy of sentences into the true and the false.

either Mars was dry at t or Mars was not dry at t . The question is composed of expressions that are not distinctively philosophical in character: “Mars,” “always,” “either . . . or . . .,” “not,” “was,” and “dry.” All of them occur in a recognizably unphilosophical question such as “Was Mars always either uninhabited or not dry?,” which someone might ask on judging that Mars is both uninhabited and dry and wondering whether there is a connection. Although philosophical issues can be raised *about* the words in both questions, it does not follow that merely in using those words one is in any way engaging in philosophy. One difference between the two questions is that it is not obviously futile to try to argue from the armchair that Mars was always either dry or not dry, whereas it is obviously futile to try to argue from the armchair that Mars was always either uninhabited or not dry.

The original question does not itself *ask* whether it is metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. It simply asks whether Mars always *was* either dry or not dry. Expressions such as “metaphysically necessary,” “knowable *a priori*,” “analytic,” and “logically true” do not occur in the original question; one can understand it without understanding any such philosophical terms of art. This is of course neither to deny nor to assert that it *is* metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. For all that has been said, the proposition may be any combination of those things. But that is not what the original question asks.

In other circumstances, we could have answered the original question on philosophically uninteresting grounds. For instance, if there had never been liquid on Mars, then it would always have been dry, and therefore either dry or not dry. In order to pose a question which could not possibly be answered in that boring way, someone who already grasped one of those philosophically distinctive concepts might ask whether it is metaphysically necessary, or knowable *a priori*, or analytic, or logically true that Mars was always either dry or not dry. The meaningfulness of the philosophical jargon might then fall under various kinds of suspicion, which would extend to the question in which it occurred. But the original question itself cannot be correctly answered in the boring way with respect to the originally envisaged circumstances. Its philosophical interest, however contingent, is actual.

We could generalize the original question in various ways. We might ask whether *everything* is always either dry or not dry. Then we might notice that discussing that question is quite similar to discussing whether everything is either old or not old, and so on. We might, therefore, ask whether for every property everything either has it or lacks it. The coherence of such generalizing over properties might itself fall under various kinds of suspicion, which would extend to the question in which it occurred. Someone might even doubt whether there is such a property as dryness. But the original question itself does not attempt such generality. That it has the same kind of philosophical interest as many other questions does not imply that it has itself no philosophical interest. If that interest is obscured by problematic features of the apparatus with which we try to generalize it, we can refrain from generalizing it, and stick with the original question. In order not to be distracted by extraneous issues that arise from the apparatus of generalization, not from the original question, we do best to stick with the original question in its concrete form.³ We can still help ourselves not to be distracted by unimportant features of the question, if we remember that there are many other questions of a similar form.

What is the original question about? “About” is not a precise term. On the most straightforward interpretation, a sentence in a context is about whatever its constituents refer to in that context. Thus, taken at face value, the original question is about the planet Mars, the referent of “Mars” in this context; perhaps it is also about dryness, the referent of “dry,” and the referents of other constituents too. Since the original question contains no metalinguistic expressions, it is not about the name “Mars” or the adjective “dry.” Evidently, the original question is not explicitly about words.

Is the original question implicitly about language? Someone might claim so on the grounds that it is equivalent to questions that are explicitly about language, such as these:

Is the sentence “Mars was always either dry or not dry” true? (Does it express a truth as used in this context?)

Did Mars always belong either to the extension of the word “dry” or to the anti-extension of “dry” (as the word “dry” is used in this context)?

³ See also Quine (1970: 11).

But parallel reasoning would lead to the conclusion that the unphilosophical question “Was Mars always either uninhabited or not dry?” is also implicitly about language, since it is equivalent to these questions:

Is the sentence “Mars was always either uninhabited or not dry” true?
(Does it express a truth as used in this context?)

Did Mars always belong either to the extension of the word “uninhabited” or to the anti-extension of “dry” (as the word “dry” is used in this context)?

Indeed, we could make parallel arguments for all everyday and scientific questions. Since they are not all about language in any distinctive sense, the reasoning does not show that the original question was about language in any distinctive sense. Even if the equivalences did show that the original question was in some sense implicitly about language, they could be read in both directions: they would also show that the explicitly metalinguistic questions were in an equally good sense implicitly not about language.

The equivalences between the questions are in any case uncontroversial only if the corresponding disquotational biconditionals are:

- (T1) “Mars was always either dry or not dry” is true if and only if Mars was always either dry or not dry.
- (T2a) For any time t , Mars belongs to the extension of “dry” at t if and only if Mars is dry at t .
- (T2b) For any time t , Mars belongs to the anti-extension of “dry” at t if and only if Mars is not dry at t .

On the face of it, these biconditionals express at best contingent truths. For perhaps the word “dry” could have meant *wet*, in which case Mars would have belonged to the extension of “dry” when wet and to the anti-extension of “dry” when dry: for *we* use the word “dry” to mean *dry* even when we are talking *about* circumstances in which it would have meant something else, because we are not talking *in* those circumstances. If so, T2a and T2b do not express necessary truths. Similarly, perhaps the sentence “Mars was always either dry or not dry” could have failed to express a truth even though Mars

was always either dry or not dry, since “always” could have meant *never*. On this reading, T1 does not express a necessary truth. We should not assume that a useful notion of aboutness would transfer across merely contingent biconditionals. Perhaps we can instead interpret T1, T2a, and T2b as expressing necessary truths by individuating linguistic expressions so that their semantic properties are essential to them; whether that requires treating the quoted expressions as necessary existents is a delicate matter. In any case, some theorists of vagueness have denied even the actual truth of biconditionals such as T1, T2a, and T2b; they might respond to the original question in one way and to the explicitly metalinguistic questions in another.⁴ Thus the questions are not pragmatically, dialectically or methodologically equivalent within the context of debates on vagueness. For present purposes, we need not resolve the status of the disquotational biconditionals, because we have already seen that the sense in which they make the original question implicitly about words is too indiscriminate to be useful.

We can argue more directly that the original question is not implicitly about the word “dry” by appeal to a translation test. For consider the translation of the original question into another language, such as Serbian:

Da li je Mars uvek bio suv ili nije bio suv?

The Serbian translation is not implicitly about the English word “dry.” But since the questions in the two languages mean the same, what they are implicitly about (in the same context) should also be the same. Therefore, the original question is not implicitly about the word “dry.” By similar reasoning, it is not about any word of English or any other language. Of course, given the informality of the notion of implicit aboutness, the argument is not fully rigorous. Nevertheless, the translation test emphasizes how far one would have to water down the notion of reference in order to reach a notion of implicit aboutness on which the original question would be implicitly about a word.

⁴ A recent example of a supervaluationist rejecting such disquotational equivalences for borderline cases is Keefe (2000: 213–20). For further discussion see Williamson (1994a: 162–4) and McGee and McLaughlin (2000).

philosophical, is not about thought or language in any distinctive sense. It does not support the linguistic or conceptual turn, interpreted as a conception of the subject matter of philosophy.

2

If the original question, read literally, had too obvious an answer, either positive or negative, that would give us reason to suspect that someone who uttered it had some other meaning in mind, to which the overt compositional structure of the question might be a poor guide. But competent speakers of English may find themselves quite unsure how to answer the question, read literally, so we have no such reason for interpreting it non-literally.

It is useful to look at some proposals and arguments from the vagueness debate, for two reasons. First, they show why the original question is hard, when taken at face value. Second, they show how semantic considerations play a central role in the attempt to answer it, even though it is not itself a semantic question.

The most straightforward reason for answering the original question positively is that “Mars was always either dry or not dry” is a logical truth, a generalization over instances of the law of excluded middle ($A \vee \neg A$, “It is either so or not so”) for various times. In my view, that reasoning is sound. However, many think otherwise. They deny the validity of excluded middle for vague terms such as “dry.”

The simplest way of opposing the law of excluded middle is to deny outright when Mars is a borderline case that it is either dry or not dry, and therefore to answer the original question in the negative. For instance, someone may hold that Mars was either dry or not dry at time t only if one can know (perhaps later) whether it was dry at t , given optimal conditions for answering the question (and no difference in the history of Mars): since one cannot know, even under such conditions, whether it is dry when the case is borderline, it is not either dry or not dry. One difficulty for this negative response to the original question is that it seems to imply that in a borderline case Mars is neither dry nor not dry: in other words, both not dry and not not dry. That is a contradiction, for “not not dry” is the negation of “not dry.”

Intuitionistic logic provides a subtler way to reject the law of excluded middle without denying any one of its instances. Intuitionists ground logic in states of increasing but incomplete information, rather than a once-for-all dichotomy of truth and falsity. They deny that anything can be both proved and refuted, but they do not assert that everything can be either proved or refuted. For intuitionists, the denial of an instance of excluded middle ($\neg(A \vee \neg A)$, “It is not either so or not so”) entails a contradiction ($\neg A \ \& \ \neg\neg A$, “It is both not so and not not so”), just as it does in classical logic, and contradictions are as bad for them as for anyone else. Thus they cannot assert that Mars was once not either dry or not dry ($\exists t \neg(\text{Dry}(m, t) \vee \neg\text{Dry}(m, t))$), for that would imply that a contradiction once obtained ($\exists t (\neg\text{Dry}(m, t) \ \& \ \neg\neg\text{Dry}(m, t))$, “Mars was once both not dry and not not dry”), which is intuitionistically inconsistent. However, although intuitionists insist that proving an existential claim in principle involves proving at least one instance of it, they allow that disproving a universal claim need not in principle involve disproving at least one instance of it. The claim that something lacks a property is intuitionistically stronger than the claim that not everything has that property. Thus one might assert that Mars was not always either dry or not dry ($\neg\forall t (\text{Dry}(m, t) \vee \neg\text{Dry}(m, t))$), on the general grounds that there is no adequate procedure for sorting all the times into the two categories, without thereby committing oneself to the inconsistent existential assertion that it was once not either dry or not dry. Hilary Putnam once proposed the application of intuitionistic logic to the problem of vagueness for closely related reasons.⁶ Thus one might use intuitionistic logic to answer the original question in the negative.

On closer inspection, this strategy looks less promising. For a paradigm borderline case is the worst case for the law of excluded middle (for a term such as “dry” for which threats to the law other than from vagueness are irrelevant), in the sense that both proponents and opponents of the law can agree that it holds in a paradigm borderline case only if it holds universally. In symbols, if Mars was a paradigm borderline case at time τ : $(\text{Dry}(m, \tau) \vee \neg\text{Dry}(m, \tau)) \rightarrow$

⁶ For intuitionist logic in general see Dummett (1977). For its application to the problem of vagueness see Graff and Williamson (2002: 473–506) and Chambers (1998).

$\forall t (\text{Dry}(m, t) \vee \neg \text{Dry}(m, t))$ (“If Mars was either dry or not dry at time τ , then Mars was always either dry or not dry”). But on this approach the law does not always hold in these cases ($\neg \forall t (\text{Dry}(m, t) \vee \neg \text{Dry}(m, t))$, “Mars was not always either dry or not dry”), from which intuitionistic logic allows us to deduce that it does not hold in the paradigm borderline case ($\neg (\text{Dry}(m, \tau) \vee \neg \text{Dry}(m, \tau))$, “Mars was not either dry or not dry at τ ”), which is a denial of a particular instance of the law, and therefore intuitionistically inconsistent (it entails $\neg \text{Dry}(m, \tau) \ \& \ \neg \neg \text{Dry}(m, \tau)$, “Mars was both not dry and not not dry at τ ”). Thus the intuitionistic denial of the universal generalization of excluded middle for a vague predicate forces one to deny that it has such paradigm borderline cases. The latter denial is hard to reconcile with experience: after all, the notion of a borderline case is usually explained by examples.

The problems for the intuitionistic approach do not end there. One can show that the denial of the conjunction of any finite number of instances of the law of excluded middle is intuitionistically inconsistent.⁷ The denial of the universal generalization of the law over a finite domain is therefore intuitionistically false too. If time is infinitely divisible, the formula $\forall t (\text{Dry}(m, t) \vee \neg \text{Dry}(m, t))$ generalizes the law over an infinite domain of moments of time, and its denial is intuitionistically consistent, but the possibility of infinitely divisible time is not crucial to the phenomena of vagueness. We could just as well have asked the original question about a long finite series of moments at one-second intervals; it would have been equally problematic. The classical sorites paradox depends on just such a finite series: a heap of sand consists of only finitely many grains, but when they are carefully removed one by one, we have no idea how to answer the question “When did there cease to be a heap?” To deny that Mars was dry or not dry at each moment in the finite series is intuitionistically inconsistent. Thus intuitionistic logic provides a poor basis for a negative answer to the original question.

Other theorists of vagueness refuse to answer the original question either positively or negatively. They refuse to assert that Mars was always either dry or not dry; they also refuse to assert that it was not always either dry or not dry.

⁷ One proves by mathematical induction on n that if \mathbf{A}_n is the conjunction of n instances of excluded middle then $\neg \mathbf{A}_n$ is intuitionistically inconsistent.

A simple version of this approach classifies vague sentences (relative to contexts) as true (T), false (F) or indefinite (I); borderline sentences are classified as indefinite. The generalized truth-tables of a three-valued logic are used to calculate which of these values to assign to a complex sentence in terms of the values assigned to its constituent sentences. The negation of A , $\neg A$, is true if A is false, false if A is true and indefinite if A is indefinite:

| | |
|----------|----------------------------|
| A | $\neg A$ |
| T | F |
| I | I |
| F | T |

A conjunction $A \& B$ (“A and B”) is true if every conjunct is true; it is false if some conjunct is false; otherwise it is indefinite. A disjunction $A \vee B$ (“Either A or B”) is true if some disjunct is true; it is false if every disjunct is false; otherwise it is indefinite:

| A | B | $A \& B$ | $A \vee B$ |
|----------|----------|--------------------------------|------------------------------|
| T | T | T | T |
| T | I | I | T |
| T | F | F | T |
| I | T | I | T |
| I | I | I | I |
| I | F | F | I |
| F | T | F | T |
| F | I | F | I |
| F | F | F | F |

A universal generalization is treated as if it were the conjunction of its instances, one for each member of the domain: it is true if every instance is true, false if some instance is false, and otherwise indefinite. An existential generalization is treated as if it were the disjunction of the instances: it is true if some instance is true, false if every instance is false, and otherwise indefinite. The three-valued tables generalize the familiar two-valued ones in the sense that one recovers the latter by deleting all lines with “I.”

Let us apply this three-valued approach to the original question. If Mars is definitely dry or definitely not dry at t (the time denoted by t), then $\text{Dry}(\mathbf{m}, t)$ is true or false, so the instance of excluded middle

$\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t)$ is true. But if Mars is neither definitely dry nor definitely not dry at t , then $\text{Dry}(\mathbf{m}, t)$ is indefinite, so $\neg\text{Dry}(\mathbf{m}, t)$ is indefinite too by the table for negation, so $\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t)$ is classified as indefinite by the table for disjunction. Since Mars was once a borderline case, the universal generalization $\forall t (\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t))$ has a mixture of true and indefinite instances; hence it is classified as indefinite. Therefore its negation $\neg\forall t (\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t))$ is also indefinite. Thus three-valued theoreticians who wish to assert only truths neither assert $\forall t (\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t))$ nor assert $\neg\forall t (\text{Dry}(\mathbf{m}, t) \vee \neg\text{Dry}(\mathbf{m}, t))$. They answer the original question neither positively nor negatively.

Three-valued logic replaces the classical dichotomy of truth and falsity by a three-way classification. Fuzzy logic goes further, replacing it by a continuum of degrees of truth between perfect truth and perfect falsity. According to proponents of fuzzy logic, vagueness should be understood in terms of this continuum of degrees of truth. For example, “It is dark” may increase continuously in degree of truth as it gradually becomes dark. On the simplest version of the approach, degrees of truth are identified with real numbers in the interval from 0 to 1, with 1 as perfect truth and 0 as perfect falsity. The semantics of fuzzy logic provides rules for calculating the degree of truth of a complex sentence in terms of the degrees of truth of its constituent sentences. For example, the degrees of truth of a sentence and of its negation sum to exactly 1; the degree of truth of a disjunction is the maximum of the degrees of truth of its disjuncts; the degree of truth of a conjunction is the minimum of the degrees of truth of its conjuncts. For fuzzy logic, although the three-valued tables above are too coarse-grained to give complete information, they still give correct results if one classifies every sentence with an intermediate degree of truth, less than the maximum and more than the minimum, as indefinite.⁸ Thus the same reasoning as before shows that fuzzy

⁸ This point does not generalize to the semantics of conditionals in fuzzy logic, given the popular rule that if the consequent is lower than the antecedent in degree of truth then the degree of truth of the conditional falls short of 1 by the amount by which the consequent falls short of the antecedent in degree of truth; otherwise the degree of truth of the conditional is 1. Hence if A has a higher degree of truth than B but both are indefinite then $A \rightarrow B$ is indefinite while $B \rightarrow A$ is perfectly true. Thus the information that the antecedent and consequent are indefinite does not determine whether the conditional is indefinite.

different ways, talking past each other. But that is not so: almost everyone who reflects on the original question finds it difficult and puzzling. Even when one has settled on an answer, one can see how intelligent and reasonable people could answer differently while understanding the meaning of the question in the same way. If it has an *obvious* answer, it is the answer “Yes” dictated by classical logic, but those of us who accept that answer can usually imagine or remember the frame of mind in which one is led to doubt it. Thus the original question, read literally, has no unproblematically obvious answer in any sense that would give us reason to suspect that someone who asked it had some other reading in mind.

Without recourse to non-literal readings, some theorists postulate ambiguity in the original question. For example, some three-valued logicians claim that “not” in English is ambiguous between the operators \neg (strong negation) and $\neg\Delta$ (weak negation): although $\neg A$ and $\neg\Delta A$ have the same value if A is true or false, $\neg\Delta A$ is true while $\neg A$ is indefinite if A is indefinite. While $A \vee \neg A$ (“It is so or not so”) can be indefinite, $A \vee \neg\Delta A$ (“It is so or not definitely so”) is always true. On this view, the original question queries $\forall t (\text{Dry}(m, t) \vee \neg\text{Dry}(m, t))$ on one reading, $\forall t (\text{Dry}(m, t) \vee \neg\Delta\text{Dry}(m, t))$ on another; the latter is true (Mars was always either dry or not definitely dry) while the former is indefinite. Thus the correct answer to the original question depends on the reading of “not.” It is “Indefinite” if “not” is read as strong negation, “Yes” if “not” is read as weak negation. Although the three-valued logician’s reasoning here is undermined by higher-order vagueness, that is not the present issue.¹⁰

If “not” were ambiguous in the way indicated, it would still not follow that the dispute over the original question is merely verbal. For even when we agree to consider it under the reading of “not” as strong negation, which does not factorize in the manner of $\neg\Delta$, we still find theories of vagueness in dispute over the correct answer. We have merely explained our terms in order to formulate more clearly a difficult question about Mars.

Still, it might be suggested, the dispute between different theories of vagueness is verbal in the sense that their rival semantics characterize different possible languages or conceptual schemes: our choice of which of them to speak or think would be pragmatic, based on

¹⁰ See Williamson (1994a: 193–5).

considerations of usefulness rather than of truth. Quine defended a similar view of alternative logics (1970: 81–6).

To make sense of the pragmatic view, suppose that the original vague atomic sentences are classifiable both according to the bivalent scheme as true or false and according to the trivalent scheme as definitely true, indefinite or definitely false, and that the truth-tables of each scheme define intelligible connectives, although the connective defined by a trivalent table should be distinguished from the similar-looking connective defined by the corresponding bivalent table. Definite truth implies truth, and definite falsity implies falsity, but indefiniteness does not discriminate between truth and falsity: although all borderline atomic sentences are indefinite, some are true and others false. As Mars dries, “Mars is dry” is first false and definitely false, then false but indefinite, then true but indefinite, and finally true and definitely true. However, this attempted reconciliation of the contrasting theories does justice to neither side. For trivalent logicians, once we know that a sentence is indefinite, there is no further question of its truth or falsity to which we do not know the answer: the category of the indefinite was introduced in order not to postulate such a mystery. Similarly, for fuzzy logicians, once we know the intermediate degree of truth of a sentence, there is no further question of its truth or falsity to which we do not know the answer: intermediate degrees of truth were introduced in order not to postulate such a mystery. In formal terms, trivalent and fuzzy logics are undoubtedly less convenient than bivalent logic; the justification for introducing them was supposed to be the inapplicability of the bivalent scheme to vague sentences. If a bivalent vague language is a genuinely possible option, then the trivalent and fuzzy accounts of vagueness are mistaken. Conversely, from a bivalent perspective, the trivalent and fuzzy semantics do not fix possible meanings for the connectives, because they do not determine truth conditions for the resultant complex sentences: for example, the trivalent table for \neg does not specify when $\neg A$ is true in the bivalent sense. It would, therefore, be a fundamental misunderstanding of the issue at stake between theories of vagueness to conceive it as one of a pragmatic choice of language.

We already speak the language of the original question; we understand those words and how they are put together; we possess the concepts they express; we grasp what is being asked. That semantic

knowledge may be necessary if we are to know the answer to the original question.¹¹ It is not sufficient, for it does not by itself put one in a position to arbitrate between conflicting theories of vagueness. For each of those theories has been endorsed by some competent speakers of English who fully grasp the question.

Competent speakers may of course fail to reflect adequately on their competence. Although the proponents of conflicting theories of vagueness presumably have reflected on their competence, their reflections may have contained mistakes. Perhaps reflection of sufficient length and depth on one's competence would lead one to the correct answer to the original question. But the capacity for such more or less philosophical reflection is not a precondition of semantic competence. Philosophers should resist the professional temptation to require all speakers to be good at philosophy.

We can distinguish two levels of reflection, the logical and the metalogical. In response to the original question, logical reflection involves reasoning with terms of the kind in which the question is phrased; the aim is to reach a conclusion that answers the question. For example, one might conclude by classical logic that Mars was always either dry or not dry; one might conclude by fuzzy logic that it is indefinite whether it was always one or the other. The logical level is not purely mechanical. When the reasoning is complex, one needs skill to select from the many permissible applications of the rules one sequence that leads to an answer to the question. When the reasoning is informal, one needs good judgment to select only moves that really are permissible applications of the rules. But one is still thinking about whatever the question was about. One starts only at the metalogical level of reflection to think about the semantics of the logical connectives and other expressions one employed at the logical level. For example, at the metalogical level one may assert or deny

¹¹ Of course, monolingual speakers of another language may know whether Mars was always dry or not dry without ever hearing of the original question, which is an interrogative sentence of English; they use a synonymous sentence of their own language. They do not know whether the original English question has a positive answer. Someone may even know whether the original English question has a positive answer without understanding the question, because the knowledge can be passed along a chain of testimony; understanding of the original question is needed only at one end of the chain. These quibbles do not affect the argument.

that the sentence “Mars was always either dry or not dry” is a logical truth. The rules used at the logical level are articulated only at the metalogical level.

It must be possible to think logically without thinking metalogically, for otherwise by the same principle thinking metalogically would involve thinking metametalogically, and so *ad infinitum*: our thinking never goes all the way up such an infinite hierarchy. What can prompt ascent to the metalogical level are hard cases in which one feels unclear about the permissibility of a given move at the logical level. One’s mastery of the language and possession of concepts leave one quite uncertain how to go on. In the case of the original question, a salient line of classical reasoning leads to a positive answer: it persuades some competent speakers while leaving others unconvinced. Even to discuss the contentious reasoning we must semantically ascend. We cannot hope to resolve the dispute undogmatically if we never leave the lower level.

3

The argument so far has reached two conclusions at first sight hard to reconcile with each other. First, the original question is not about thought or language. Second, to answer it adequately one must assess rival theories of vagueness in thought and language. How can that way of reaching an answer be appropriate to the original question? We might, therefore, find ourselves tempted back to the idea that somehow the original question was surreptitiously about thought or language.

On further reflection, the combination of the two conclusions is less surprising. Many non-philosophical questions that are not about thought or language cannot be resolved without inquiry into thought or language. Suppose that a court of law must decide whether Smith killed Jones. The question is not who said or thought what. Nevertheless, the crucial arguments may be over whether to trust the witnesses’ testimony. How is what they say now related to what they think now or thought then? How is what they think now or thought then related to what actually happened? Are they lying or sincere? Are their memories confused or clear? Those are questions about their thought and speech. They hold the key to whether Smith killed Jones, even

though that question is not about thought about language.¹² Of course, the questions about the thought or talk are not about it in isolation from what it is thought or talk about: they are relevant because they concern the relation between the thought or talk and what it is about.

The court must decide the issue on the evidence before it. In a criminal case, does the evidence put it beyond reasonable doubt that Smith killed Jones? In a civil case, does the evidence make it more probable than not? If the court is really deciding a question about testimonial evidence, that is already a question about talk.¹³ But the question about the evidence arises in virtue of its bearing on the primary question, whether Smith killed Jones. Indeed, the question about the evidence is exactly a question about its bearing on the primary question. So the point stands.

Historians are often in a similar position. They want to know what happened. The way to achieve that is largely by considering documents, linguistic accounts of what happened – not in isolation, but in relation to what they represent. Most obviously, historians want to know whether the documents accurately represent what happened, but to answer that question they must in turn ask about their provenance: who produced them, when and why? Thus the history of the events of primary interest requires a history of thought and talk about those events. Those histories typically overlap, for thought or talk about some part of a complex human event is often another part of the same complex event.

Something analogous occurs in the methodology of the natural sciences. We wish to know the value of some physical quantity. We must devise apparatus to measure it. We may find ourselves in disputes over the functioning of different devices. Although the primary

¹² The issue of Smith's intentions concerns his thoughts, but we may suppose that the question immediately at issue is whether Smith was even involved in Jones's death.

¹³ Non-testimonial evidence may be taken to include non-linguistic items such as a bloodied knife; this is what lawyers call "real evidence." For an argument that all evidence in an epistemologically central sense of the term is propositional see Williamson (2000a: 194–200). For example, the evidence in this sense might include the proposition that the bloodied knife was found at the scene of the crime, but not the knife itself.

technical, metalinguistic, and metaconceptual theories. This phenomenon seems to have been overlooked by those who complain about the “arid” technical minuteness of much philosophy in the analytic tradition. A question may be easy to ask but hard to answer. Even if it is posed in dramatic and accessible terms, the reflections needed to select rationally between rival answers may be less dramatic and accessible. Such contrasts are commonplace in other disciplines; it would have been amazing if they had not occurred in philosophy. Impatience with the long haul of technical reflection is a form of shallowness, often thinly disguised by histrionic advocacy of depth. Serious philosophy is always likely to bore those with short attention-spans.¹⁴

Why should considerations about thought and language play so much more central a role in philosophy than in other disciplines, when the question explicitly under debate is not itself even implicitly about thought or language? The paradigms of philosophical questions are those that seem best addressed by armchair considerations less formal than mathematical proofs. The validity of such informal arguments depends on the structure of the natural language sentences in which they are at least partly formulated, or on the structure of the underlying thoughts. That structure is often hard to discern. We cannot just follow our instincts in reasoning; they are too often wrong (see Chapter 4 for details). In order to reason accurately in informal terms, we must focus on our reasoning as presented in thought or language, to double-check it, and the results are often controversial. Thus questions about the structure of thought and language become central to the debate, even when it is not primarily a debate about thought or language.

The rise of modern logic from Frege onwards has provided philosophers with conceptual instruments of unprecedented power and precision, enabling them to formulate hypotheses with more clarity and determine their consequences with more reliability than ever before. Russell’s theory of descriptions showed vividly how differences between the surface form of a sentence and its underlying semantic structure might mislead us as to its logical relations and thereby create philosophical illusions. The development of formal

¹⁴ Popularization has its place, in philosophy as in physics, but should not be confused with the primary activity.

model-theory and truth-conditional semantics by Tarski and others has provided a rigorous framework for thinking about the validity of our inferences. These theoretical advances have enormous intellectual interest in their own right. They may have made it tempting to suppose that all philosophical problems are problems of language: but they do not really provide serious evidence for that conjecture.

To deny that all philosophical questions are about thought or language is not to deny the obvious, that many are. We have also seen how in practice the attempt to answer a question which is not about thought or language can largely consist in thinking about thought and language. Some contemporary metaphysicians appear to believe that they can safely ignore formal semantics and the philosophy of language because their interest is in a largely extra-mental reality. They resemble an astronomer who thinks he can safely ignore the physics of telescopes because his interest is in the extra-terrestrial universe. In delicate matters, his attitude makes him all the more likely to project features of his telescope confusedly onto the stars beyond. Similarly, the metaphysicians who most disdain language are the most likely to be its victims. Again, those who neglect logic in order to derive philosophical results from natural science make frequent logical errors in their derivations; their philosophical conclusions do not follow from their scientific premises. For example, some supposed tensions between folk theory and contemporary science depend on fallacies committed in the attempt to draw out the consequences of common sense beliefs.

Analytic philosophy at its best uses logical rigor and semantic sophistication to achieve a sharpness of philosophical vision unobtainable by other means. To sacrifice those gains would be to choose blurred vision. Fortunately, one can do more with good vision than look at eyes.

Many have been attracted to the idea that all philosophical problems are linguistic or conceptual through the question: if the method of philosophy is *a priori* reflection, how can it lead to substantive knowledge of the world? Those who find that question compelling may propose that it informs us of relations of ideas rather than matters of fact, or that its truths are analytic rather than synthetic, or that it presents rules of grammar disguised as descriptions, or that its aim is the analysis of thought or language. In short, on this view, philosophical truths are conceptual truths. We may suspect the pres-

ence of empiricist presuppositions in the background – or, as with Ayer, in the foreground. Not starting with such presuppositions, we should be open to the idea that thinking just as much as perceiving is a way of learning how things are. Even if one does not fully understand *how* thinking can provide new knowledge, the cases of logic and mathematics constitute overwhelming evidence that it does so. The case of the original question, which is philosophical yet queries a theorem of classical logic, shows that we cannot segregate logic from philosophy and claim that armchair thinking illuminates the former but not the latter. In particular, conceptions of logic and mathematics as (unlike philosophy) somehow trivial or non-substantial have not been vindicated by any clear explanation of the relevant sense of “trivial” or “non-substantial.” Whether a given formal system of logic or mathematics is consistent is itself a non-trivial question of logic or mathematics. We know from Gödel’s second incompleteness theorem that the consistency of most standard systems of elementary mathematics cannot be decided in equally elementary mathematics, unless the original system is already inconsistent. The next two chapters investigate in more depth the prospects for conceptual truth and its role in philosophy.

3

Metaphysical Conceptions of Analyticity

1

“Philosophical questions are more conceptual in nature than those of other disciplines”: that can easily pass for a statement of the obvious.¹ Many philosophers consciously seek conceptual connections, conceptual necessities, conceptual truths, conceptual analyses. In effect, they present themselves as seeking far more general and less obvious analogues of “Vixens are female foxes.” The suggestion is that an armchair methodology is appropriate to their quest because it concerns truths in some sense less substantial, less world-involving than those of other disciplines: in Humean terms, relations of ideas rather than matters of fact. Our conceptual or linguistic competence, retained in the armchair, is to suffice for *a priori* knowledge of the relevant truths.

As already argued, philosophical truths are not generally truths *about* words or concepts. However, analytic truths are not supposed to be always about words or concepts, even if words or concepts are supposed to play a special role in explaining their truth. The sentence “Vixens are female foxes” is in no useful sense about the word

¹ To give just one example, even Jack Smart, whose work robustly engages the nature of the non-linguistic, non-conceptual world and who described metaphysics as “a search for the most plausible theory of the whole universe, as it is considered in the light of total science” (1984: 138), could also write that philosophy is “in some sense a *conceptual* inquiry, and so a science can be thought of as bordering on philosophy to the extent to which it raises within itself problems of a conceptual nature” (1987: 25), although he admits that he “cannot give a *clear* account of what I have meant when earlier in this essay I have said that some subjects are more concerned with “conceptual matters” than are others” (1987: 32).

“vixen” or any other words; it is about vixens, if anything. Its meaning is not to be confused with that of the metalinguistic sentence “‘Vixens are female foxes’ is true.” Similarly, the thought *vixens are female foxes* is not about the concept *vixen* or any other concepts; it too is about vixens, if anything. It is not to be confused with the metaconceptual thought *the thought VIXENS ARE FEMALE FOXES is true*.

How can a sentence which comes as close as “Vixens are female foxes” does to being a definition of “vixen” be about vixens rather than about the word “vixen”? Uttering it in response to the question “What does ‘vixen’ mean?” normally enables the questioner to work out the answer to the question, by pragmatic reasoning, even though the literal meaning of the sentence does not directly answer the question, just as does uttering “That is a gnu” while pointing at one in answer to the question “What does ‘gnu’ mean?.” If core philosophical truths are analytic, they may *exhibit* significant features of words or concepts without *describing* them.

Does the conception of philosophical truths as analytic or conceptual vindicate a form of the linguistic or conceptual turn without misrepresenting the subject matter of philosophy as itself linguistic or conceptual? The case study in the previous chapter gave no support to such a conjecture. Nevertheless, let us examine the matter more systematically.

Many philosophically relevant truths are clearly not conceptual truths in any useful sense. For instance, in arguing against subjective idealism, a defender of common sense metaphysics says that there was a solar system millions of years before there was sentient life. Similarly, a defender of common sense epistemology says that he knows that he has hands; that he knows that he has hands is no conceptual truth, for it is consistent with all conceptual truths that he lost them in a nasty accident. Some philosophers of time argue that not only the present exists by appeal to Special Relativity. Philosophers of mind and language dispute whether there is a language of thought; whatever the answer, it is no conceptual truth. Naturalists and anti-naturalists dispute whether there is only what there is in space and time; again, the answer is unlikely to be a conceptual truth. Moral and political philosophers and philosophers of art appeal to empirically discovered human cognitive limitations, and so on. Such philosophical arguments cannot be dismissed on general

interchangeable and “analytic” was taken to do the work of both. But that does not yet imply that no work remains for it to do.

If we try to sort sentences as “analytic” or “synthetic” in the manner of chicken-sexers, we can usually achieve a rough consensus. Of course borderline cases will occur, but so they do for virtually every distinction worth making: perfect precision is an unreasonable demand. The issue is what theoretical significance, if any, attaches to the rough boundary thus drawn. Even if “analytic” is defined in terms of “synonymous” and other expressions under better control than “analytic,” we should not assume without checking that it has any of the consequences sometimes associated with it. In particular, we should not assume that analytic truths are insubstantial in any further sense.

Nothing in this book challenges the legitimacy of familiar semantic terms such as “synonymous.” They will be used without apology, and they permit various senses of “analytic” to be defined. But none of them makes sense of the idea that analytic truths are less substantial than synthetic ones, or that core philosophical truths are less substantial than the truths of most other disciplines. There is something robust about “Two Dogmas of Empiricism”: insights remain even when its skepticism towards meaning is stripped away.

On some conceptions, analytic sentences are true simply in virtue of their meaning, and analytic thoughts simply in virtue of their constituent concepts. They impose no constraint on the world, not even on that part of it which consists of words and concepts. That is why it is unnecessary to get up out of one’s armchair to investigate whether such a constraint is met. Analytic truths are less substantial than synthetic ones because the latter do impose constraints on the world, which it may or may not meet. This is another way of putting the idea that analytic truths are true in virtue of meaning alone while synthetic truths are true in virtue of a combination of meaning and fact, for if analytic truths did impose constraints on the world, they would be true partly in virtue of the fact that the world met those constraints, and so not true in virtue of meaning alone. Call such conceptions of analyticity *metaphysical*. Other conceptions dispense with the idea of truth in virtue of meaning, and treat analyticity as a privileged status in respect of knowledge or justification which a sentence or thought has in virtue of the conditions for understanding its constituent words or possessing its constituent concepts. Although the privileged truths

impose constraints on the world, the task of checking that they are met is somehow less substantial than for other truths, for those who understand the relevant words or possess the relevant concepts. Call such conceptions of analyticity *epistemological*.⁴

This chapter examines a variety of attempts to develop a metaphysical account of analyticity. Some depend on misconceptions about meaning or truth. Others yield intelligible notions of analyticity, by watering down the traditional account to a point where it loses most of its usually supposed implications. They provide no reason to regard analytic truths as in any way insubstantial.⁵ Even if core philosophical truths are analytic in such a sense, that does not explain how we can know or justifiably believe them.⁶ At best it reduces the problem to the epistemology of another class of truths, such as necessary truths or logical truths. The next chapter will examine attempts to develop an epistemological account of analyticity, also with negative results. The overall upshot is that philosophical truths are analytic at most in senses too weak to be of much explanatory value or to justify conceiving contemporary philosophy in terms of a linguistic or conceptual turn.

The conclusion is not best put by calling purportedly analytic truths “substantial,” because in this context the term “substantial” is hopelessly vague. Rather, appeals in epistemology to a metaphysical conception of analyticity tend to rely on a *picture* of analytic truths as imposing no genuine constraint on the world, in order to

⁴ See Boghossian (1997) for the distinction between metaphysical and epistemological accounts of analyticity, and Tappenden (1993: 240) for a somewhat similar distinction.

⁵ Etchemendy (1990: 107–24) contrasts “substantive” generalizations with logical ones. The idea is widespread. It occurs in different forms in Wittgenstein’s *Tractatus Logico-Philosophicus* and in Locke’s “Of trifling propositions” (*An Essay Concerning Human Understanding*, Book IV, Chapter viii).

⁶ Since analytic truths are standardly taken to be sentences, the term “true” will sometimes be applied to sentences, as well as to thoughts and propositions; where required, the context makes clear what kind of truth-bearer is intended. Talk of knowing or believing a sentence should be understood as elliptical for talk of having knowledge or belief which one can express with the sentence (on its standard meaning). Thus someone who knows “Grass is green” knows that grass is green and can express that knowledge by saying “Grass is green”; this is not to be confused with the metalinguistic knowledge that the sentence “Grass is green” is true.

explain the supposed fact that knowing them poses no serious cognitive challenge. If that account could be made good, it would provide a useful sense for “insubstantial,” which would refer to the pictured property, epistemological not in its nature but in its explanatory power. Substantial truths would be the ones that lacked this property. But the account cannot be made good. The metaphysical picture cannot be filled in so as to have the required explanatory power in epistemology. Thus “substantial” and “insubstantial” are not provided with useful senses. The negation of a picture is not itself a picture. That is a problem for appeals to metaphysical analyticity, not for the present critique.

2

The distinction between analytic truth and synthetic truth does not distinguish different *senses* of “true”: analytic and synthetic truths are true in the very same sense of “true.” That should be obvious. Nevertheless, it is hard to reconcile with what many logical positivists, Wittgensteinians and others have said about analytic truths. For they have described them as stipulations, implicit definitions (partial or complete), disguised rules of grammar and the like. On such a conception, enunciating an analytic truth is not stating a fact but something more like fixing or recalling a notation: even if talk of truth as correspondence to the facts is metaphorical, it is a bad metaphor for analytic truth in a way in which it is not for synthetic truth. In the face of this conception, we should remind ourselves why “truth” is quite unequivocal between “analytic truth” and “synthetic truth.”

We can start by considering a standard disquotational principle for truth (where both occurrences of “P” are to be replaced by a declarative sentence):

(T) “P” is true if and only if P.

If “true” is ambiguous between analytic truth and synthetic truth, (T) must itself be disambiguated. Nevertheless, the left-to-right direction holds for both notions:

- (Talr) “P” is analytically true only if P.
 (Tslr) “P” is synthetically true only if P.

Obviously, “Bachelors are unmarried” is analytically true only if bachelors are unmarried, just as “Bachelors are untidy” is synthetically true only if bachelors are untidy. The exact parallelism of (Talr) and (Tslr) already casts doubt on the supposed ambiguity. Indeed, they are jointly equivalent to a single principle about the disjunction of analytic truth and synthetic truth (“simple truth”):

- (Taslr) “P” is analytically true or synthetically true only if P.

Worse, the right-to-left direction fails for both notions:

- (Tarl) “P” is analytically true if P.
 (Tslr) “P” is synthetically true if P.

For (Tarl) has a false instance when a synthetic truth is substituted for “P”; (Tslr) has a false instance when an analytic truth is substituted for “P.” There are no natural substitutes for the right-to-left direction of (T) in the form of separate principles for analytic truth and synthetic truth. Rather, the natural substitute for the right-to-left direction disjoins the two notions:

- (Tasrl) “P” is analytically true or synthetically true if P.

But (Tasrl) reinstates simple truth as the theoretically important characteristic.

One cannot avoid the problem by qualifying “true” in (T) with “analytic” for “the relevant kind of sentence” and with “synthetic” for the rest. For the sentences of the relevant kind are presumably just the analytic truths and analytic falsehoods. Thus the schemas for analytic and synthetic truth amount to these:

- (Ta) If “P” is analytically true or analytically false, then “P” is analytically true if and only if P.
 (Ts) If “P” is neither analytically true nor analytically false, then “P” is synthetically true if and only if P.

But (Ta) and (Ts) follow from (Taslr), (Tasrl) and the analogue for falsity of (Taslr):⁷

(Faslr) “P” is analytically false or synthetically false only if not P.

Thus the information in (Ta) and (Ts) is in effect just information about the disjunction of analytic truth and synthetic truth. The attempt to treat analytic truth and synthetic truth separately just confuses the theory of “true.” The same happens for other theoretically important applications of “true.”

Consider the standard two-valued truth-table for the material conditional:

| A | B | A → B |
|---|---|-------|
| T | T | T |
| T | F | F |
| F | T | T |
| F | F | T |

If “true” is ambiguous between analytic truth and synthetic truth, what does “T” mean in that table? We might try subscripting it as T_{analytic} and $T_{\text{synthetic}}$, multiplying the possibilities in the first two columns accordingly and adding the appropriate subscript in the third column. “F” will require corresponding subscripts too. Since the possibilities T_{analytic} , $T_{\text{synthetic}}$, F_{analytic} and $F_{\text{synthetic}}$ arise for both A and B, the new truth-table will have sixteen lines. Worse, consider this case:

⁷ Proof: Assume (Taslr), (Faslr) and (Tasrl). To derive (Ta), note that it is equivalent to the conjunction of two claims: (i) if “P” is analytically true, then “P” is analytically true if and only if P; (ii) if “P” is analytically false, then “P” is analytically true if and only if P. Now (i) is logically equivalent to the claim that “P” is analytically true only if P, which follows from (Taslr). Moreover, by (Faslr) “P” is analytically false only if not P; as just seen “P” is analytically true only if P, so “P” is analytically false only if “P” is not analytically true; thus if “P” is analytically false then both sides of the biconditional in the consequent of (ii) fail, so (ii) holds. To derive (Ts), first note that “P” is synthetically true only if P by (Taslr). Conversely, if P then “P” is analytically true or synthetically true by (Tasrl); since by the antecedent of (Ts) it is not analytically true, it is synthetically true. Incidentally, by themselves (Ta) and (Ts) are weak in other ways too; in particular, they do not entail that nothing can be both analytically true and synthetically true.

lawyers. It is no good to say “Never mind whether barristers *are* lawyers; ‘Barristers are lawyers’ is true simply because it means that barristers are lawyers.” For any true sentence *s* whatsoever, a canonical explanation of the truth of *s* takes the overall form “*s* means that *P*, and *P*.”⁹ To use the obscure locution “in virtue of,” every true sentence is true in virtue of both its meaning and how things are. This is another way of making the point that analytic truths and synthetic truths are not true in radically different ways.¹⁰

We can ask “in virtue of” questions about non-metalinguistic matters too. In virtue of what are vixens female foxes? To use another obscure locution, what makes it the case that vixens are female foxes? An appeal to semantic or other facts about the words “vixen,” “female” and “fox” in answer to those questions would confuse use and mention. Vixens would have been female foxes no matter how we had used words. Presumably, vixens are female foxes in virtue of whatever female foxes are female foxes in virtue of; what makes it the case that vixens are female foxes is whatever makes it the case that female foxes are female foxes. Some may argue that female foxes are not female foxes in virtue of anything; nothing makes it the case that female foxes are female foxes. The suggestion may be that analytic truths require no truthmaker, unlike synthetic truths. An alternative suggestion is that analytic truths require truthmakers of a different kind from those of synthetic truths. Such suggestions are too unconstrained to be tractable for assessment. Still, two points stand out. First, they seem to conflict with general principles of

⁹ See Boghossian (1997: 335–6). Quine says that we can say that the logical truth “Everything is self-identical” depends for its truth “on an obvious trait, viz., self-identity, of its subject matter, viz., everything.” However, he claims that it makes no difference whether we say that or say that it depends for its truth “on traits of the language (specifically on the usage of “=”), and not on traits of its subject matter” (1966: 106).

¹⁰ Another problem for the supposed contrast is that it seems to equivocate on “means.” When we explain why “Barbara is a barrister” is true by saying “It means that Barbara is a barrister, and Barbara *is* a barrister,” “means” can be paraphrased as “expresses the proposition”; what proposition a sentence expresses may depend on the context in which it is uttered, if indexicals are present. By contrast, the appeal to meaning in the case of analytically true sentences is not to the proposition expressed on some particular occasion but rather to the linguistic meaning of the sentence, which is invariant across contexts, even if indexicals are present.

truthmaker theory (in the unlikely event that such a theory is needed). For instance, what makes a disjunction true is what makes one of its disjuncts true. Thus whatever makes (2) (“Barbara is a lawyer”) true also makes both (5) and (6) true:

- (5) Barbara is a lawyer or Barbara is not a lawyer.
- (6) Barbara is a lawyer or Barbara is a doctor.

But (5) is a simple logical truth, while (6) is a straightforward synthetic truth. Second, no connection has been provided between truthmaker theory and epistemology. Knowing a truth need not involve knowing its truthmaker; one can know (6) without knowing which disjunct is true (Barbara works in a building where only lawyers and doctors work). No account has been given as to why it should be easy from an armchair to know a truth with no truthmaker, or a truthmaker only of the special sort supposedly appropriate to analytic truths.

Nevertheless, at least one clear difference between paradigms of “analytic” and paradigms of “synthetic” is in the vicinity. For meaning that barristers are lawyers is sufficient for being true, whereas meaning that Barbara is a barrister is not. More generally, call a meaning *sufficient for truth* just in case necessarily, in any context any sentence with that meaning is true.¹¹ Thus the meaning of “Barristers are lawyers” is sufficient for truth; the meaning of “Barbara is a barrister” is not. One proposal is to explicate “analytic truth” as “truth whose meaning is sufficient for truth.” Call this “modal-analyticity.”¹² For non-skeptics about meaning and necessity, the

¹¹ To handle ambiguity, treat it as homonymy: distinct sentences with the same superficial form. The reification of meanings in the definition can be eliminated at the cost of circumlocution. Note also that the utterance of a modal-analytic truth may be false if the context shifts during the utterance: consider “If it is now exactly noon then it is now exactly noon.” Similarly, an utterance of “If John is a bachelor then John is unmarried” may express a falsehood if the wedding ceremony is completed between the utterance of the antecedent and the utterance of the consequent. Taking such complications into account would not help friends of analyticity.

¹² The notion of modal-analyticity is similar to the notion of deep necessity in Evans (1979), where the truth of the sentence does not depend on any contingent feature of reality.

notion of modal-analyticity is quite intelligible. But what are its consequences?

Consider any non-indexical sentence s that expresses a necessarily true proposition. Necessarily, in any context, any sentence with the actual meaning of s expresses that necessary truth and is therefore true. Thus s is a modal-analytic truth, because its meaning is sufficient for truth. In that sense, it is true in virtue of meaning. But how little has been achieved in so classifying it! Nothing has been done to rule out the hypothesis that it expresses a profound metaphysical necessity about the nature of the world, knowable if at all only through arduous *a posteriori* investigation, for instance. No reason has been provided to regard s as “merely verbal” or “insubstantial” in a pretheoretic sense, unless one already had independent reason to regard all necessities as merely verbal or insubstantial. Similarly, mathematical truths count as modal-analytic; their so counting is by itself no reason to regard them as merely verbal or insubstantial. Indeed, for all that has been said, even “Water contains H_2O ” is modal-analytic, given that “water” has a different meaning as used on Twin Earth to refer to XYZ, a different substance with the same superficial appearance.

To make the point vivid, call a meaning *temporally sufficient for truth* just in case at all times, in any context any sentence with that meaning is true. Read the quantifiers “at all times” and “in any context” non-modally, so they do not range outside the actual world. Thus any sentence which expresses, in a time-independent way, an eternally true proposition, however contingent, has a meaning temporally sufficient for truth. For example, the meaning of “No hotel ever has a billion rooms” is presumably temporally sufficient for truth. We can call the sentence “temporal-analytic” if we like, but that in no way implies that it is somehow insubstantial, because there is no background connection between eternity and some sort of insubstantiality. Similarly, calling a sentence “analytic” in the sense of modal-analyticity does not imply that it is somehow insubstantial, in the absence of a background connection between necessity and some sort of insubstantiality. Yet the account of analyticity was what was supposed to substantiate the claim of insubstantiality. If we already had a background connection between necessity and insubstantiality, there would be little to gain from invoking modal-analyticity in order to argue that core philosophical truths are insubstantial, since

we could do it more simply just by arguing that true philosophical sentences in the core express necessarily true propositions.

Admittedly, not all modal-analytic true sentences express necessarily true propositions. Examples of the contingent *a priori* such as “It is raining if and only if it is actually raining” are modal-analytic, since the truth of “It is raining” as uttered in a given context is necessarily equivalent to the truth of “It is actually raining” as uttered in that context, because “actually” refers rigidly to the world of the context, but the biconditional does not express a necessary truth, since the weather could have been relevantly different, in which case it would have been not raining if and only if it is actually raining. Thus modal-analyticity violates Kripke’s constraint that analyticity implies necessity; in this respect it may diverge from the traditional conception. Conversely, not all sentences that express necessarily true propositions are modal-analytic: consider examples of the necessary *a posteriori* such as “I am not Tony Blair.” Nevertheless, such examples seem marginal to the envisaged conception of core philosophical truths, most of which will both express necessarily true propositions and be modal-analytic.

A core of philosophical truths may indeed be modal-analytic. Some philosophers seek to articulate necessary truths without essential reliance on indexicals; if they succeed, the sentences they produce are modal-analytic. Even if contextualists are right, and key philosophical terms such as “know” shift their reference across contexts, the relevant sentences may still both express necessarily true propositions and be modal-analytic: consider “Whatever is known to be the case is the case.” The answers to philosophical questions of the forms “Is it possible that P?” and “Is it necessary that P?” will themselves express necessary truths, given the principle of the widely accepted modal logic S5 that the possible is non-contingently possible and the necessary non-contingently necessary; if the answers can be phrased in non-indexical terms, they will then be modal-analytic. But outside the envisaged core many philosophically relevant truths will not be modal-analytic, as the examples near the start of the chapter show.

Unfortunately, even for modal-analytic philosophical truths, classifying them as modal-analytic does not unlock their epistemology, any more than classifying a truth as necessary explains how we can know it. Of course, if a sentence is modal-analytic, then one is safe from error in uttering it with its given meaning. In that sense, one’s utterance is reliable. But such reliability falls well short of what

knowledge requires, since otherwise any true mathematical assertion would count as an expression of knowledge, no matter how fallacious the “proof” on which it was based. “Vixens are female foxes” is utterly misleading as a paradigm for the epistemology of modal-analytic truths in general. To say that *s* is a modal-analytic truth whose constituent words and grammar we understand does very little way to explain how we can know or justifiably believe *s*.¹³ In particular, it does not imply that the mere linguistic understanding of *s*, which every competent speaker possesses, provides any insight into the truth of *s*, or constitutes more than the minimal starting-point for inquiry it does for ordinary synthetic truths.

4

Issues related to those just raised for modal-analyticity arise for what is sometimes called “Frege-analyticity.”¹⁴ A sentence is Frege-analytic just in case it is synonymous with a logical truth. For example, “All furze is furze” is a logical truth, roughly speaking because everything of the form “All F is F” is true. “All furze is gorse” is not a logical truth, because not everything of the form “All F is G” is true (“All fungus is grease” is false). However, “All furze is gorse” is Frege-analytic, because it is synonymous with the logical truth “All furze is furze,” since “furze” is synonymous with “gorse.” In “Two Dogmas,” Quine admits the notion of logical truth, and therefore allows that if “synonymous” were legitimate, so would be “analytic” in the sense of Frege-analyticity. By present standards, the notion of Frege-analyticity is quite intelligible. But what are its consequences?

Trivially, every logical truth is Frege-analytic, because it is synonymous with itself. Clearly, this alone does nothing to show that logical truths are somehow insubstantial in any metaphysical, epistemologically explanatory sense (see the end of Section 1). For instance, it is compatible with the hypothesis that there are truths of second-order logic which characterize the necessary structure of reality in profound

¹³ See n. 6 for this terminology.

¹⁴ The term “Frege-analytic” is from Boghossian (1997), with reference to §3 of Frege (1950) (as Boghossian suggests, the interpretation of the passage is not entirely clear). He classifies the notion of Frege-analyticity as neither epistemological nor metaphysical but semantic (1997: 363); for convenience, it is treated here under the heading of metaphysical notions of analyticity.