



truth as one and many

MICHAEL P. LYNCH

OXFORD

Truth as One and Many

This page intentionally left blank

Truth as One and Many

Michael P. Lynch

CLARENDON PRESS · OXFORD

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

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

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

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

Published in the United States
by Oxford University Press Inc., New York

© Michael P. Lynch 2009

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

First published 2009

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

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

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloguing in Publication Data

Lynch, Michael P. (Michael Patrick), 1966–
Truth as one and many / Michael P. Lynch.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-19-921873-8

1. Truth. I. Title.

BD171.L8695 2009

121—dc22

2008053080

Typeset by Laserwords Private Limited, Chennai, India

Printed in Great Britain
on acid-free paper by
The MPG Books Group

ISBN 978-0-19-921873-8

10 9 8 7 6 5 4 3 2 1

For Kathleen

This page intentionally left blank

Contents

<i>Preface</i>	ix
Introduction	i
1. Truisms	7
2. Truth as One	21
3. Truth as Many	51
4. Truth as One and Many	69
5. Truth, Consequence, and the Universality of Reason	85
6. Deflationism and Explanation	105
7. Expanding the View: Semantic Functionalism	129
8. Applying the View: Moral Truth	159
Conclusion	191
<i>Select Bibliography</i>	193
<i>Index</i>	201

This page intentionally left blank

Preface

I became interested—some would say obsessed—with the issue of truth early on in my career as a philosopher. The present book presents my current thinking on the subject. Its central thesis is that truth can manifest itself in more than one way in our cognitive life. The book articulates this view, defends it, and then extends and applies it.

So many people have influenced the development of the ideas presented here, either through written comments or conversation, that the following can at best constitute only a partial list: Robert Barnard, Donald Baxter, Jc Beall, Tom Bontly, David Capps, Marian David, Michael Devitt, Tim Elder, Pascal Engel, Hartry Field, Chris Gauker, Patrick Greenough, Steven Hales, Eberhard Herrmann, Joel Kupperman, Matt McGrath, Daniel Massey, Philip Pettit, Tom Polger, Mark Richard, Marcus Rossberg, Stewart Shapiro, Gila Sher, Robbie Williams, and Elia Zardini. Conversations with Crispin Wright in particular over the past decade have been a constant source of inspiration; while we often disagree about the details, we agree on the big issues, and I always learn something from him. Others who have similarly influenced—and provoked—my thinking about truth include William Alston, Simon Blackburn, Terry Horgan and Paul Horwich.

Early versions of core chapters were presented at various universities, including the University of Cincinnati, University of Nancy, University of Genoa, University of Stirling, Uppsala University, Tufts University, Bloomsburg University, among others. In 2005 and 2006, several chapters were given an airing in seminars at Arché Centre for Philosophical Research in Language, Logic, Metaphysics and Epistemology at the University of St Andrews, from which I benefited immensely. A version of the manuscript was read in a seminar I gave at UCONN in 2007; comments

from the seminar's participants led to many improvements; of particular help were Colin Caret, Alexis McCloud, Ian Smith and Aaron Cotnoir. Cotnoir's extensive, page-by-page comments in particular saved me from many errors. Lionel Shapiro likewise improved the manuscript with comments on core chapters, as did comments on a still later version from Nic Damnjanovic, Doug Edwards, Nikolaj Pedersen, Adam Podlaskowski, Cory Wright and two anonymous reviewers for Oxford University Press. Paul Bloomfield, as always, supplied general encouragement as well as extremely helpful comments, particularly on the final chapter. I am especially grateful for the support of my editor, Peter Momtchiloff, and to Javier Kalhat for preparing the index. Thanks to one and all; naturally, all errors are mine alone.

Work on the final versions of the manuscript was supported by a Fellowship at the University of Connecticut Humanities Center in 2006–2007, and a University of Connecticut Provost's Fellowship in the spring of 2008. Portions of some chapters have appeared as “Alethic Pluralism, Logical Consequence and the Universality of Reason”, *Midwest Studies in Philosophy*, XXXII (2008), 122–140; “Truth and Multiple Realizability”, *Australasian Journal of Philosophy*, 82 (2004), 384–408; “A Coherent Moral Relativism” (with D. Capps and D. Massey) *Synthese*, in press, “ReWrighting Pluralism”, *The Monist*, 89 (2006), 63–85.

My deepest debt is to my wife and muse Terry Berthelot, and to my daughter Kathleen. Their love is the light that has led me through this and many other follies.

August 2008
Mansfield CT

Introduction

1. Unity, Diversity and a Puzzle

You and I believe all kinds of propositions—that torture is morally barbaric, that the tallest tree in the front yard is a spruce, that twice sixteen is thirty-two. Yet despite obvious differences in subject matter, we ordinarily assume that these kinds of propositions—ranging from morality to the sundries of everyday life—are equally capable of being correct or incorrect, that they can get things right, or get them wrong. In short, we intuitively treat them all as if they can be true. Yet when we look around the world for the objects and properties that some of these beliefs purportedly represent, we find ourselves at a loss. With regard to at least some of the propositions we believe, there appears to be no objects, properties, facts—in short, no reality—to which they correspond.

To take just one example, consider the proposition that it is morally wrong to torture another human being. This certainly seems true. But it is puzzling *how* it can be true if it must correspond to mind-independent reality to do so. For while it is beyond debate that people are sometimes tortured, it is highly debatable that there are objective moral values. This is particularly so if naturalism is our theoretical background. It is difficult to know how to “locate” something like moral wrongness amongst the furniture of the physical world. Moral wrongness—and hence the fact that torture has that property—doesn’t seem to be the sort of thing that we can naturally investigate or discern, nor does it seem to be the sort of thing that is physical or supervenes on the physical. What it is, and how we know about it, can seem mysterious.

Of course, not everyone will be gripped by this mystery in the case of moral propositions. For any given subject matter, one

can always resist the idea that there is anything puzzling about how propositions about that subject correspond to reality. This is so wherever we find *realism* plausible—wherever we think there are mind-independent objects and properties that our beliefs are representing.

But few will be confidently realist across the board. Most of us will continue to find it puzzling how we can apparently believe certain kinds of propositions and yet not be able to explain how they could correspond to some reality. The question is what to say about that.

The usual options have come to seem—at least to me—somewhat tired. They are alike in accepting that truths must correspond to some reality, alike in declaring that some troublesome propositions do not do so, and alike in denying that said propositions are literally true. They differ chiefly in the intuitions they privilege. Some, such as expressivists, are impressed by the *diversity* of our thought, by the different functions that our various thought contents play in our lives. They conclude that some thoughts serve as vehicles of sentiment rather than representation; they are therefore neither true nor false, but in a different game entirely. Others, such as error theorists, emphasize instead the *cognitive unity* of our thought contents—the fact that despite their radical differences in subject matter, all our judgments and beliefs seem equally apt for rational assessment. Moral beliefs, for example, *aim* to represent reality, but according to this account, they just fail to do so.

Even partisans must acknowledge that both intuitions—about diversity and unity—are pre-theoretically appealing. So it shouldn't be surprising that it has proven difficult to privilege one at the expense of the other. The content of our thought as we've found it is both diverse *and* unified—both open to being true and yet radically different in subject and function. The trick, I've come to believe, is how to make sense of this in the face of our puzzle.

In my view, we can begin to make sense of it only by rejecting an assumption that gets the puzzle going: namely, that if a belief or

its content is true, it must be true in the same way—for example, by corresponding to reality.

This book develops and defends a new theory of truth that rejects just this assumption. According to what I'll call the functionalist theory, truth is a functional property that can be realized—or, as I shall say in the book “manifested”—in more than one way. Theories, especially new theories, need motivation. And one motivation for adopting the functionalist theory of truth is that if it were correct, we would have at the least the start of a new and more satisfying solution—or dissolution—to our puzzle. For if truth is a functional property, then our true beliefs about the concrete physical world needn't manifest truth—or “be true in the same way”—as our thoughts about matters where the human stain is deepest, such as morality or the law. Consequently, if the theory developed here is correct, then it is at least possible that some of the propositions we believe may be true without having to correspond to reality, or represent objects in the world. If so, then the way is open to understanding how our various thought contents can be both diverse in kind and yet cognitively unified.

2. A Sketch of the Territory

A second motivation for adopting a functionalist theory of truth is more direct. It has benefits that other theories of truth lack.

Many contemporary philosophers—like most philosophers over the course of Western philosophical history—are *monists* about truth; they assume that there is one and only one explanation of what makes something true. Like gold or potassium, they think that truth has a single inner structural essence—a philosophical “atomic number”. Of course, they disagree over what truth's nature actually is, whether it is a matter of correspondence between thought and world or a type of idealized coherence among our beliefs. But they agree that where a proposition we believe is true, it is true in the same way.

In recent decades, many philosophers have come to think that the monist's quest for the nature of truth is a fool's errand. A commonly cited reason is that monist theories all seem open to devastating counterexamples. They face what I call in the book the scope problem: for any sufficiently robustly characterized truth property *F*, there appear to be some kind of proposition *K* that lack *F* but that are intuitively true (or capable of being true).

The new orthodoxy is some version or other of deflationism. Rather than signaling a special property that all and only true propositions have in common, the deflationist takes it that the concept of truth is a mere expressive device, useful for purposes of generalization and semantic ascent. Truth, or rather "true", is an honorific that all propositions therefore compete for equally.

The simplicity of the deflationist picture can be appealing. And it seems at first blush to suggest an easy solution to the puzzle noted above. For unlike the traditional expressivist or error theorist, the deflationist can happily accept that moral judgments are true or false. And since there is no special property in which being true consists, there is no special problem of trying to figure out whether all kinds of propositions can have it. Ascribing truth to the judgment that torture is wrong is no more or less informative, no more or less objective, no more or less mysterious, than the judgment itself. Moreover, it seems we could still acknowledge, with the traditional expressivists, that moral judgments function to express attitudes. If deflationism were right, they would still be capable of being true or false in the same deflationary sense that every judgment is capable of being true or false. And this seems like a happy result: we would then seem to get diversity and unity too.

But the benefits of deflationism come with significant costs. As we'll see in greater detail later, chief among them is that deflationism *removes truth from our explanatory toolkit*. And that means that we must relinquish the most obvious explanation of diversity. For consider now what a view such as the expressivism cum deflationism just envisioned can say about the difference in content between our moral and non-moral judgments. As

deflationists are well aware, such a view cannot account for that difference by appealing to truth. We cannot say that what makes the content of moral judgments or beliefs distinct from the content of beliefs about the physical world is that the latter but not the former have “objective truth-conditions” or “correspond to fact”. To do so would be tantamount to rejecting deflationism. We would be appealing to the nature of truth to explain why some judgments differ in content from others. If the deflationist is right, truth has no nature. Consequently, we cannot appeal to it to help explain other items of philosophical interest such as content. And I think that this should give us pause. These latter items are difficult enough to understand without barring ourselves in advance from appealing to some of the more obvious tools at our disposal.

The alternative view I defend in this book is at odds with both traditional monists and their deflationary critics, but it also has something in common with both views. Deflationists are right to be skeptical of the thought that any one traditional theory of truth can tell us what all and only true beliefs have in common. At a suitable level of abstraction, understanding what true beliefs are involves simply understanding what they *do*—their role in our cognitive economy. To play this role is to satisfy certain truisms, truisms that display truth’s connections to other concepts. It is this truth-role that gives truth its unity; the features that are constitutive of this role are what true propositions have in common, and simply having those features is what we ordinarily mean by saying that a proposition is true. But not all facts about truth are exhausted by the truisms. One such fact is that there is more than one property that can make beliefs true. Truth, as I’ll put it, is *immanent* in these other properties of beliefs. In some domains, what makes a belief true is that it corresponds to reality; in others, beliefs are made true by a form of coherence. Traditional theories are therefore right to insist in the face of their deflationary critics that there is more to say about truth, and that what more there is to say can help us explain other items of interest: like the diversity of content. But they are wrong in that the traditional theories are not best conceived of

as theories of truth itself. They are better seen as theories of the properties that make beliefs true—or *manifest truth*.

In sum, the view I'll be defending can be thought of as having two components. The first is a functionalist analysis of both the ordinary concept of truth and the property that concept is a concept of. The second is the thesis that this one property can be manifested in more than one way. Thus the overall position incorporates a form of *pluralism* about truth. But it is not a simple ambiguity view of truth; it does not imply that "true" simply has different meanings when appended to different beliefs. Truth is immanent in *distinct* properties of beliefs; our ordinary concept of truth is univocal.

The broad motivations I've just cited for taking the functionalist view seriously are best appreciated, of course, in light of a full discussion of that theory. I've organized that discussion as follows. Chapter 1 addresses the obvious—if often neglected—question of what makes a theory of truth a theory of truth. Chapters 2 and 3 collectively make the case that we should be neither traditional monists nor simple pluralists about truth. Chapters 4 and 5 articulate the functionalist theory and defend it against objections. Chapter 6 distinguishes it from its more deflationary rivals. Chapter 7 offers a tentative expansion of functionalism to other key semantic concepts. Chapter 8 offers an application of the view to the difficult case of moral truth.

The theory—it might better be called a picture—of truth that emerges in these pages is distinctive and, I hope, clear. But it is far from comprehensive. The questions addressed are chiefly metaphysical ones, chief among them the nature of truth. Some issues of profound importance, particularly formal issues regarding the semantic paradoxes, are regretfully left untouched. I make no apology for this. Trying to get a full picture of truth is like trying to get a full picture of the world; it is only possible from very far away.

1

Truisms

If a person shows that such things as wood, stones, and the like, being many are also one, we admit that he shows the coexistence of the one and the many, but he does not show that the many are one or the one many; he is uttering not a paradox but a truism.

Plato, *Parmenides*

I. Truisms about Truth

My question is Pilate's: what is truth? But unlike Pilate, I aim to take this question seriously and answer it head-on. But before we can do so, we first need to consider a preliminary question. What would *count* as a theory of the nature of truth? By this I mean, how do we know whether some theory is about *truth* as opposed to being about some other thing?

In metaphysics we aim to find the nature of things and properties that puzzle us—be it pain, or causation, or identity or truth. But we can't search for that which we know nothing about. So when setting off to discover the nature of some target property it helps to have some understanding of what it is we are looking for: its *nominal essence*, as Locke might have put it. The nominal essence of F, in the sense I intend, is our folk concept of F. It embodies our preconceptions, the way we tacitly think about it in ordinary life—even if, normally, we don't even recognize ourselves as doing so. A natural way of identifying something's nominal essence,

therefore, is to appeal to the set of largely implicit beliefs we folk have about it. By appealing to those folk beliefs, or truisms, we won't typically learn *everything* about the object or property we are interested in. And our later discoveries may force us to revise our preconceptions of it—especially when the something in question is natural, like gold, or water or magnetism. At the very least, our later substantive theories of the property may help us to see that some of our folk beliefs about it are more important and central than others. But however these questions play out, keeping one eye on our folk beliefs about the thing about which we are curious will hopefully tell us whether our subsequent theories of its nature address the topic we were concerned with when our theorizing began.

So before setting off to discuss various theories of the nature of truth, let's briefly consider a few of our folk beliefs about it. I will try to state these preconceptions as intuitively as I can, passing over for the moment various technicalities.

One preconception most of us share is that truth is objective. To speak truly is to “say of what is, that it is”, as Aristotle famously put it.¹ And since what we say, at least when we are sincere, is an expression of what we believe or judge, a parallel truism holds about belief. That is,

Objectivity: The belief that p is true if, and only if, with respect to the belief that p, things are as they are believed to be.

The truth of a belief depends on how things are; not on how I or anyone else might wish them to be.² Believing, as we say, doesn't make it so.

Objectivity is a central truism about truth. Together with some further and reasonably obvious assumptions, it underwrites further derivative principles which are typically highlighted by philosophers. One related principle is that when, for example, I

¹ *Metaphysics* Γ. 7.27, trans. Christopher Kirwan (Oxford: Oxford University Press, 1993).

² Compare W. P. Alston, *A Realist Conception of Truth* (Ithaca, NY: Cornell University Press, 1996), 22 ff and Künne, *Conceptions of Truth* (Oxford: Oxford University Press, 2003), 333 ff.

believe that roses are red, things are as I believe them to be just when roses are red. That is,

With respect to the belief that *p*, things are as they are believed to be if, and only if, *p*.

With this point in hand, we can derive, together with Objectivity, instances of:

BS: The belief that *p* is true if and only if *p*.

Another related thought is that what is true when my belief or judgment is true is the content of my judgment or belief. Thus when I believe that roses are red, strictly speaking it is not my act of believing that is true but *what* I believe, namely that roses are red. If, following philosophical convention, we call that which I believe or disbelieve a *proposition*, then we can further derive

TS: The proposition that *p* is true if, and only if, *p*.

BS and TS are the doxastic and propositional versions of the T-schema; the philosophers' favorite truism about truth. As we'll have occasion to remark upon later, many take TS, in particular, to be *the* central principle about truth. But often little or nothing is said about why such a principle—with its more theoretical commitment to propositions—should be found so compelling. Our line of reasoning suggests an explanation: TS is a natural consequence of the Objectivity truism, together with certain obvious facts about belief's relation to truth.

In committing ourselves to the idea that truth is objective, we commit ourselves to the twin hallmarks of Objectivity: the possibility of error and ignorance. What we believe to be so may not be what is, and what is we may not believe to be so. What holds for belief also holds for warranted belief. That is, we are apt to reject that every instance of the following must be true

(w) *p* if and only if the belief that *p* is warranted.

If one had lived in the tenth century, one might have been warranted in believing that the earth is flat when it is not. And

there are certainly some propositions, such as the propositions that it rained on this spot 15,000 years ago, or that the number of stars in the universe now is even, for which we lack evidence, either for them or for their negation.³ They are “undecidable”. But that hardly entails that they are not true. Consequently, if we accept BS, but *reject* (w), we accept

Warrant Independence: Some beliefs can be true but not warranted and some can be warranted without being true.⁴

As these examples indicate, the Objectivity truism underwrites several other key truisms about truth. Indeed, that truth is objective is often thought by philosophers to be our most fundamental preconception about it. And perhaps it is. But to focus on it exclusively would be to forget that truth is not only objective; it is also valuable. This fact reveals itself in two other truisms about truth. The first is the thought that, as William James put it, truth is “the good in the way of belief”.⁵ That is,

Norm of Belief: It is prima facie correct to believe that p if and only if the proposition that p is true.⁶

This is a truism, but it is not trivial; the left-hand side doesn’t merely restate the right. What is true is the propositional content of the belief, while what is correct is the believing of that content. Thus the two sides of Norm of Belief state different facts; while Norm of Belief *as a whole* claims those facts are co-extensive.

³ So here, at least, we can agree with Rumsfeld’s dictum “the absence of evidence is not evidence for absence”.

⁴ Compare Alston, *A Realist Conception of Truth*; C. Wright, *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992) 20–1; and Wright, “Minimalism, Deflationism, Pragmatism, Pluralism”, in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA: MIT Press, 2001), 751–88.

⁵ W. James, *Pragmatism and the Meaning of Truth* (Cambridge, MA: Harvard University Press, 1942).

⁶ Some may object to the biconditional, preferring instead to understand the norm as stating a necessary condition for correct belief. See P. Engel, *Truth* (London: Acumen Press, 2002): for further discussion of the present formulation, see M. P. Lynch, *True to Life* (Cambridge, MA: MIT Press, 2004). As noted below, the functionalist theory itself can survive disagreement over the best way to state the truisms.

Norm of Belief tells us that truth is belief's basic norm of correctness. It is widely held that this fact is part of what distinguishes believing from various other cognitive attitudes.⁷ Imagining, assuming, and hoping, for example, are each governed by norms—assumptions can be justified or not, imaginings can be sharp or vague, hopes can be rational or irrational. But unlike believing, neither imagining that *p*, assuming that *p*, nor hoping that *p* is properly evaluated in terms of truth. Moreover, belief is indirectly *responsive* to truth. In the typical conscious, deliberative case, belief is indirectly responsive to truth by being directly responsive to evidence. It is correct to believe what is based on evidence *because* beliefs based on evidence are likely to be true.

So Norm of Belief is an important fact about belief. As David Velleman has put it, “for a propositional attitude to be a belief just is, in part, for it to be capable of going right or wrong by being true or false”.⁸ Yet the fact that truth is the norm of belief is not just a fact about belief. It is also a truism about truth, and for the same reason that the aim of a game is to win is not just a fact about games, it is also a fact about winning.⁹ Just as the Objectivity truism connects truth with Objectivity, so Norm of Belief connects truth with the concepts of belief.

If truth is the normative standard of belief then presumably it plays a regulative role for any practice that aims at producing beliefs. Inquiry is just such a practice, and hence, not surprisingly, a third truism is that truth plays a regulative role for epistemic inquiry. Truth—or more accurately, true belief—is a goal of inquiry, as it is typically put.

⁷ See P. Boghossian, “The Normativity of Content”, *Philosophical Issues*, 13 (2003), 31–45; M. P. Lynch, “The Values of Truth and the Truth of Values”, in D. Pritchard (ed.), *Epistemic Value* (Oxford: Oxford University Press); N. Shah, “How Truth Governs Belief”, *Philosophical Review*, 112 (2003), 447–83; R. Wedgwood, “The Aim of Belief”, *Philosophical Perspectives*, 36 (2002), 267–297.

⁸ D. Velleman, *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000), 16.

⁹ See M. Dummett, “Truth”, in his *Truth and Other Enigmas* (Cambridge, MA: Harvard University Press, 1978), 1–19.

Like the link between truth and Objectivity, and truth and belief, the connection between truth and inquiry has often been highlighted by philosophers, most famously by Charles Peirce, who simply reduced truth to the aim of inquiry or to “the opinion which is fated to be ultimately agreed to by all who investigate”.¹⁰ But one needn’t go so far as Peirce to see the obvious relation between inquiry and truth. Nor must one have a specialized notion of inquiry (as Peirce may well have). By “inquiry” I mean simply the process of asking and answering questions, from the sublime “Can something come from nothing?” to the mundane (“Where are my car keys?”). Truth—in the sense of true beliefs and judgments—is clearly a goal of this process: unless the situation is highly atypical, when I ask you where my car keys are I want to know where they are—I want the truth. In pursuing inquiry of course, we pursue truth only indirectly by explicitly pursuing reasons and evidence. But we care about giving reasons, supplying justification for our beliefs, because beliefs which are so justified are more likely to be true, even if they aren’t guaranteed to be such. And this fact explains why, when we don’t know what is true, we steer by the evidence, even if evidence sometimes steers us wrong.

Of course, we don’t always pursue the truth, indirectly or otherwise. And sometimes, believing what is true isn’t the best thing—some falsehoods might be better to believe in certain circumstances and some trivial or dangerous truths may not be worth pursuing all things considered. But these cases are the exceptions that prove the rule: *other things being equal*, true beliefs are worth pursuing.

End of Inquiry: Other things being equal, true beliefs are a worthy goal of inquiry.

¹⁰ C. Peirce, “How to Make our Ideas Clear”, *Popular Science Monthly*, 12 (1878), 286–302.

In other words, it is not only correct to believe a given true proposition, other things being equal, the state of affairs of believing true propositions is worth striving for.¹¹

2. Truisms and Theory

Our folk preconceptions about truth do not—at least not obviously—drag in their wake any particular theory of what makes them true. One can implicitly recognize the link between Objectivity and truth without knowing anything about metaphysics, correspondence, “states of affairs”, or the like. Likewise with End of Inquiry and Norm of Belief: one can grant that truth is an aim of the process of asking and answering questions without having any particular view about *why* it is an aim. Those are further questions.¹²

Moreover, we should allow that some truisms—and therefore the features and relations to other properties picked out by those truisms—may well be more heavily weighted epistemically speaking than others. Call such truisms *core truisms*. Core truisms about truth cannot be denied without significant theoretical consequence and loss of plausibility. If you do deny any one of them, you must be prepared to explain how this can be so in the face of intuitive opposition. And denying *many or all* would mean that you would be regarded by other users of the concept as changing the subject.

The three historically prominent folk truisms cited above—*Objectivity*; *Norm of Belief* (and the closely associated) *End of Inquiry*—are prime candidates for core truism status. Collectively, they connect truth to the intimately related concepts of inquiry, belief, and in the case of what is arguably the most central truism,

¹¹ Here I remain neutral on whether true beliefs are the only proper end of inquiry. For present purposes we can also remain neutral on thorny issues about how best to characterize the truth goal; see Lynch, *True to Life*, and “Replies to Critics”, *Philosophical Books*, 46 (2005), 331–42.

¹² See Lynch, *True to Life*, chs. 8–10.

to objective reality—how things are. It is difficult to deny that truth has these relations in the platitudinous sense identified by the truisms. Someone, for example, who sincerely says that he believes truly that roses are red even when that is not how things are is either incoherent or not talking about the same property we are talking about when we talk about truth. Likewise, with *Norm of Belief*: someone who says that it is not even prima facie right to believe what is true is using “truth” (and probably “belief”) to talk about something other than what the rest of us use those words to talk about. The same holds, plausibly, for *End of Inquiry*.¹³

Although they are the most historically influential, these three aren’t the only plausible candidates for being core truisms. Several other candidates follow more or less directly from the historically prominent trio and some obvious premises. TS, which some see as a distillation of Objectivity, is the most obvious example. Other principles, perhaps slightly less central to our network of intuitive beliefs about truth, arguably still rank as core truisms.

One such truism we have noted previously is Warrant Independence. Another, which we employed above to derive TS from Objectivity, is

Content: It is what we believe or say that is true or false.

Content is consistent with holding that propositions are the objects of beliefs, and thus that it is propositions that are true or false. So Content by itself obviously doesn’t determine all questions about what bears truth. One might, for example, hold that propositions are true or false, but also hold that the only propositions that are capable of being true or false are propositions which are capable of being believed. Likewise one might think that propositional-truth is derivative from sentence-truth: one might think that the only propositions that can be true or false are those expressible by the sentences of natural language. And of course one might simply

¹³ For further discussion of this point about inquiry, see Lynch, “The Values of Truth and the Truth of Values”, in Haddock, Millar and Pritchard, *Epistemic Value*, forthcoming.

deny that what we say or believe are propositions in the first place. But these are additional matters that go beyond principles derivable from our truisms.

There are still other plausible candidates for core truisms. Objectivity, via TS, also grounds the further thought that:

Transparency: Whatever attitude we take towards some proposition, we are committed to taking towards its truth.

So if I believe that Bush is a lame duck President, I am also committed to believing that it is true; and if I doubt that proposition, I also doubt its truth. Like Norm of Belief, Transparency connects truth to propositional attitudes. Principles like Warrant Independence and others such as

Only true propositions can be known

illustrate the connection between truth and epistemic concepts. Still other platitudes connect truth to logical properties. Thus for example, we endorse that

Truth is what is preserved in valid inference.

Still others with moral principles, such as

True propositions are what honest people typically intend to assert.

The core truisms about truth are “core” because, as noted, denying them threatens to change the subject, or at the very least, comes at a significant theoretical cost. Of course, the core truisms don’t exhaust our folk beliefs about truth, even the folk beliefs that are universally or near universally shared. For example, many people believe that

It is very difficult to know what truth is.

But this is not the sort of belief that I call a core truism about truth. One can easily deny this principle (many philosophers do deny it) without changing the subject or accepting a deep theoretical

consequence. Nonetheless, the belief is widely shared and is worth calling a folk belief about truth.

It is worth emphasizing that one can grant that there are core truisms about truth and still hold that there can be some debate amongst philosophers about *which* principles those are. The fact that the folk have a—largely implicit—conception of some property does not imply that there will be *universal* agreement amongst theoreticians as to how best to characterize or capture that conception. Thus in order to accept the general picture, whether a principle counts as a core truism needn't be settled (nor, given the amount of folk beliefs we have about truth, can we reasonably expect it to be always settled).¹⁴ Thus, the pluralist will expect there to be substantive philosophical debate over whether

Bivalence: Every proposition is either true or false,

will be considered as a core truism about truth. Many, no doubt, will take it to be. Even more will consider the following to be core:

Non-contradiction: No proposition can be both true and false.

But of course not everyone will: those swayed by intuitions to the effect that truth can be understood as epistemically constrained in some domains may not accept the first¹⁵; those sympathetic to dialetheism may reject the second.¹⁶ For purposes of this book, I will remain agnostic about their status as members (or not) of our core truisms about truth, while emphasizing that others may disagree and still embrace the functionalist theory advocated here. Note moreover, that such agnosticism is consistent with it turning

¹⁴ This makes truth no different than most items of philosophical interest, where often there is significant debate about which beliefs about that item are central enough to help constitute our concept of that item, and which are more peripheral.

¹⁵ For discussion of this point, see for example, Wright, *Truth and Objectivity*.

¹⁶ For recent discussion of the truth and status of the law of non-contradiction, see Jc Beall, G. Priest, and B. Armour-Garb (eds.), *New Essays on Non-Contradiction* (Oxford: Oxford University Press, 2004).

out that both principles are part of our folk conception of truth without being core truisms.

Just as the approach is consistent with disagreement amongst specialists about which of our folk beliefs about truth count as core truisms, it is no embarrassment to this approach towards our folk beliefs that the above *statements* of our truisms about truth, even our statements of the core truisms, may not be immediately recognized as such by the folk.¹⁷ While some statements of some truisms—like Objectivity—will be widely regarded as truisms properly so-called, some principles that compose our folk theory may be far from platitudes in the ordinary sense of that term. But this point, while correct, needn't worry us. This is because, first, and as many others have long advocated, we can say that some of the principles and beliefs that comprise our folk theories of properties like causation or truth are believed but believed *tacitly*.¹⁸ That is, many of the folk *would* endorse platitudes like Norm of Belief or Warrant Independence were the matter to ever come up, and were they to possess the technical vocabulary that we have used to state these points.

We are now in a position to return to the question with which we began this chapter. What makes a theory a theory of truth? We can give a two-part answer.

A theory counts as a theory of *truth* (as opposed to a theory of something else) only if it incorporates the core truisms about truth. As noted, there may be disagreement amongst philosophers about just what those core truisms are. But in this book, I will take them to include, at the least, the truisms that truths are objective, correct to believe and an end of inquiry. To incorporate a truism into a theory is to either list it among the principles of the theory or endorse a principle that entails it.

¹⁷ This point bears on some of Cory D. Wright's remarks in his "On the Functionalization of Pluralist Approaches to Truth", *Synthese*, 145 (2005), 1–28. See also my reply, "Alethic Functionalism and Our Folk Theory of Truth", *Synthese*, 145 (2005), 29–43.

¹⁸ F. Jackson, *From Metaphysics to Ethics* (Oxford: Oxford University Press, 1998); D. Lewis "Mad Pain and Martian Pain", in N. Block (ed.), *Readings in Philosophy of Psychology* vol. 1 (Cambridge, MA.: Harvard University Press, 1980).

Yet to count as a *theory* of truth a theory must do more than incorporate core truisms. This is because theories *explain* the nominal essence of that which they are theories of. Hence theoretical accounts of truth, in addition to incorporating the core truisms, must *explain them* or, in the case of those that they do not incorporate, *explain them away*.

To explain the truisms, in the sense of “explain” relevant here, is to show why they are true by pointing to some property or properties that all true propositions have that results in those propositions satisfying the truisms. Traditional theories of truth such as the correspondence theory typically attempt to do just this. They attempt to explain at least most of the nominal essence of truth in terms of an underlying real essence. The thought is that correspondence is the property that all true beliefs have in common. And it is the having of this property that explains why true beliefs satisfy the central truisms—why they are objective, correct, a worthy aim of inquiry and so on.

A theory explains a truism away, on the other hand, by employing one of two strategies. First, it might supply reasons for *revising* a given truism or folk belief about truth, perhaps by arguing that the truism, as presently endorsed, isn’t universally true of all truths. Second, the theory might explain away a truism by demonstrating that, appearances to the contrary, the truism isn’t actually about *truth*. It is about something else, or perhaps nothing at all. As we’ll see, deflationary theories of truth tend to adopt this second strategy. They argue that every fact about truth can be explained by just one truism—TS. Consequently, they say that the other truisms can either be deduced from TS or are not really truisms about truth at all. Moreover, they see no reason to appeal to some further property of true propositions to explain why we accept TS. In their view, there is no need for a substantive theory of truth that serves as a foundation for our truisms. There is no need for a real essence of truth.

The view I'll defend in this book suggests an altogether different stance from either the traditional or deflationary approaches. On this view, we need the truisms to tell us what truth is, but we need the substantive theories to tell us how truth is manifested in the different domains of our cognitive life.

This page intentionally left blank

2

Truth as One

To copy reality, is, indeed, one very important way of agreeing with it, but it is far from being essential.

William James *Pragmatism*

I. Correspondence and Representation

I aim to defend the idea that beliefs can be true in different ways. In order to make sense of this idea, two things need doing. First, some reasons must be given for being unhappy with the view that there is only one way for beliefs to be true. And second, something needs to be said about these different “ways” a belief can be true: what are they exactly? We can accomplish both tasks by coming to grips with the nature and limits of the two most important traditional theories of truth.

The most venerable theory of truth is the idea that beliefs are true when they correspond to reality. Such views are often labeled *realist*, in that they start from the Objectivity truism, a truism which they are often accused of simply re-stating. Indeed, one of the most persistent objections to the correspondence theory is that it is vacuous—a mere platitude that any other theory will accept.¹ The vacuity objection is well-taken when the theory is stated as above,

¹ S. Blackburn, *Spreading the Word* (Oxford: Oxford University Press, 1984); for an extensive and important discussion of this point, see G. Vision, *Veritas* (Cambridge, MA: MIT Press, 2004). For an excellent overview and discussion of correspondence theories in general see M. David, “The Correspondence Theory of Truth”, *Stanford Encyclopedia of Philosophy* (2005): <http://plato.stanford.edu/entries/truth-correspondence/>.

or even when gussied up as, e.g. “a belief is true if and only if it corresponds to a fact”. For absent a theory of what “corresponds” means, or “facts” are, it is difficult to see how such a “theory” differs from the Objectivity truism.

Whatever the fate of the traditional versions of the correspondence theory, their descendents are not mere truisms; they involve substantive and controversial philosophical commitments. Direct successors of the correspondence theory are widely accepted within philosophy and implicitly accepted by many cognitive scientists and psychologists. Those who accept such views typically don’t see themselves as working on truth, however, but on the nature of representation.

The over-arching research program of cognitive science takes it that the mind—that is, the brain—is an organ part of whose function is to represent the world around it, so the organism whose mind/brain it is can more successfully negotiate that world. Most adherents to this program take it that the most basic representational mental items are beliefs. And beliefs not only represent, they misrepresent. Therefore a recognized goal of contemporary cognitive science is to explain what it is for beliefs to *correctly* represent. Since correctly representing beliefs are correct beliefs, and, as we noted in the last chapter, it is a truism that correct beliefs are true beliefs, a theory of what it is for mental representations to be correct counts by our criteria as a theory of truth. Call this a representational theory of truth. Representational theories are the most plausible successor to the traditional correspondence theory. Indeed, some classical correspondence theories can be understood, with little alteration, as forms of representationalism. Theories such as those held by Wittgenstein or Russell, for example, can be taken as holding that (simple, non-compound) beliefs represent, or picture, *facts*.²

² L. Wittgenstein, *Tractatus Logico-Philosophicus* (London: Routledge & Kegan Paul, 1922); B. Russell, “On the Nature of Truth and Falsehood” reprinted in his *Philosophical Essays* (London: George Allen & Unwin, 1966).

Contemporary representationists tend to be wary of positing facts as metaphysically distinct entities over and above objects and properties. There are good reasons for this. Facts are either constituted by objects and properties (and relations) or they are not. If they are, then for reasons of ontological parsimony, we must be given a serious motivation for taking them to be distinct entities over and above that which composes them. It is difficult to see what really compelling motivation could be supplied. If they are not constituted by objects and their properties, then what is their nature exactly? As Strawson famously remarked, they begin to look suspiciously like the mere shadows of statements.³

Consequently, most contemporary representationists—at least those motivated by a commitment to naturalism—unpack the correspondence metaphor slightly differently. Rather than taking beliefs to directly represent proposition-shaped entities, the thought is to take the components of those beliefs—concepts—to represent objects and properties.

Many of the core elements of what I'm calling the representational theory of truth were initially developed to understand how *sentences* and their component words represent, or *refer* to the world. But the basic elements can, and have been adapted to mental representations, to beliefs and their component concepts. And whether it is applied to sentences or beliefs, contemporary naturalistic representationalism can be understood as offering a two-part theory of truth.⁴ First, the truth of a belief, say, is defined in terms of the representational features of its component concepts

³ P. F. Strawson, "Truth", *Proceedings of the Aristotelian Society*, suppl. vol. 24 (1950), 129–56.

⁴ For early statements of the view, see H. Field, "Tarski's Theory of Truth", *Journal of Philosophy*, 69 (1972), 347–75; and M. Devitt, *Realism and Truth* (Princeton, NJ: Princeton University Press, 1984). An important recent formulation of this sort of approach can be found in T. Horgan, "Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence", in M. Lynch (ed.), *The Nature of Truth* (Cambridge, MA: MIT Press, 2001), 67–96. See also T. Horgan and R. Barnard, "Truth as Mediated Correspondence", *The Monist*, 89 (2006), 28–49. Horgan's approach, however, is contextualist and allows for different kinds of correspondence.

(what I will here call “denotation”). Thus in the case of a belief whose content has the simple predicational structure *a is F*, we get:

REPRESENT: The belief that *a is F* is true if and only if the object denoted by $\langle a \rangle$ has the property denoted by $\langle F \rangle$.⁵

The basic thought is that beliefs are true because their components stand in certain representational relations to reality and that reality is a certain way. Adopting machinery made familiar with Tarski, the representationalist then applies this insight to beliefs with more complicated structures.⁶ The result is a view according to which the truth of complex beliefs is recursively defined in terms of the truth of simpler beliefs and the rules for logical connectives, while less complex beliefs “correspond to reality” in the sense that their component parts—concepts—themselves represent objects and properties.

The second part of any representational view of truth is a theory of how concepts denote objects and properties. Of course, one might accept that truth can be defined in terms of denotation without seeing the need to say what it is. Some philosophers take the concept of denotation or reference to be explanatorily trivial, in the sense that all there is to say about it is to be found in the (infinite instances) of a schema such as

$\langle c \rangle$ denotes *x* if and only if $c = x$.⁷

Thus $\langle \text{dog} \rangle$ denotes an object just when it is a dog etc. There is nothing more to say about the matter than that. In particular, there is nothing more to be said about *why* claims of this sort should be true—why our mental representations denote the objects they do denote.

Contemporary representationalists, on the other hand, see the nature of denotation as a substantive question, one which, in

⁵ Throughout, I use brackets in the usual way: $\langle \text{dog} \rangle$ means the concept of a dog; $\langle \text{snow is white} \rangle$ means the proposition that snow is white.

⁶ A. Tarski, “The Concept of Truth in Formalized Languages”, in his *Logic, Semantics Metamathematics*, 2nd edn., trans. J. H. Woodger (Indianapolis, IN: Hackett, 1983).

⁷ See e.g. T. Horwich, *Truth*, 2nd edn. (Oxford: Oxford University Press, 1998).

principle at least, is open to naturalistic explanation. Here the representationalist is apt to regard believing and conceiving as analogous to perceiving. Perceiving is a broadly psychological process, which produces representations of objects around us. Perception, conceived of in this way, can be naturalistically investigated. It is not explanatorily trivial; it is robust relation we bear to the world around us, which is or supervenes on even more basic causal processes. Therefore, if, as REPRESENT suggests, we take our beliefs to also represent the world around us in virtue of how their component concepts refer to objects and express properties, it is a reasonable—if controversial—hypothesis that how our beliefs represent the world can also be naturalistically investigated.

The task of the contemporary representationalist is to come up with a theory that would allow us to frame such investigations. Toy versions of two familiar versions are these. First, there are broadly causal theories. Such theories come in a variety of forms. Very broadly speaking, however, we could say that for such theories,

CAUSAL: <cat> denotes cats = cats, cause, under appropriate conditions, mental tokenings of <cat>.⁸

As is well known, the plausibility of causal theories of this sort hinge on carving out a reasonable account of when conditions are “appropriate”. For while it seems plausible enough that my concept of cat wouldn’t represent cats unless, in many cases, the presence of cats caused my system to token that concept, there seems to be too many things other than cats that also cause me to token <cat>—pictures of cats, small dogs seen at a distance and so on. Consequently, some way of ruling out these possibilities—of declaring them non-standard so to speak—must be given. Doing

⁸ One could also say, for a predicative concept like <red>

CAUSAL (pred) <red> denotes redness = instances of redness, cause, under appropriate conditions, mental tokenings of <redness>.

so has been, and in some circles still is, the subject of a thriving industry.⁹

Skepticism about the likely success of this project has led some representationalists in a very different direction. One plausible idea is that <cat> represents cats and not small dogs because that concept is *supposed to* map cats. That is its purpose, what it was evolutionarily designed to do. Concepts represent what they do because that is their biological function.¹⁰ Thus to stay with our example, a simplified version of this might be:

TELEOLOGICAL: <cat> denotes cats = the biological function of <cat> is to be mentally tokened in presence of cats.

Both of these proposals require development. Both face objections. I am not interested in defending their details here. Rather, I want to stress simply that both can be thought of as a *framing hypothesis* for *naturalistically* investigating mental representation in the manner just suggested. This means that to some extent, we should not be surprised, or necessarily disheartened, when either hypothesis encounters difficulties. Naturalistic investigation often requires us to refine our framing hypothesis in the face of new information.

For our purposes, the real promise of a naturalistic theory of representation is that theories like CAUSAL and TELEOLOGICAL can be combined with REPRESENT to give a representational

⁹ Obviously I am simplifying in the text the complexities of these theories, and passing over numerous differences in formulation. For theories of this (broad) sort, see F. Dretske, *Knowledge and the Flow of Information* (Cambridge, MA: MIT Press, 1981); D. W. Stampe, "Toward a Causal Theory of Linguistic Representation", in P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in Philosophy 2* (Minneapolis, MN: University of Minnesota Press, 1977); J. Fodor, *Psychosemantics* (Cambridge, MA: MIT Press, 1987); M. Devitt, *Designation* (New York: Columbia University Press, 1981) and *Realism and Truth*, 2nd edn. (Princeton, NJ: Princeton University Press, 1997).

¹⁰ Again, I pass over numerous details and complexities irrelevant for our purposes, as well as important differences between certain views, the most important of which is how one understands "biological purpose" or "proper function". For details of teleofunctional views on this matter and others see R. Millikan, *Language, Thought and other Biological Categories* (Cambridge, MA: MIT Press, 1984); D. Papineau, *Reality and Representation* (Oxford: Blackwell, 1987); for a good overview of the differences between such views, see G. Macdonald and D. Papineau (eds.), *Teleosemantics* (Oxford: Oxford University Press, 2006).

theory of truth. According to this theory, truth is defined in terms of representation, representation is defined in terms of denotation, and denotation is defined as a property that either is, or supervenes on natural relations like those specified in CAUSAL or TELEOLOGICAL. Thus, to give another toy example, a representational view might be constructed as follows. Let's say that an object or property, which, under appropriate conditions, causes (or its instances cause) mental tokenings of some concept to be "causally mapped" by that concept. If so, we can construct:

CC (Causal-Correspondence): The belief that a is F is true if and only if the object causally mapped by $\langle a \rangle$ has the property causally mapped by $\langle F \rangle$.

Likewise with a teleological theory of representation: Let us say that a concept that has as its biological function to be mentally tokened in the presence of a particular object or property *functionally maps* that object or property. If so, then one might construct:

TC (Teleological Correspondence): The belief that a is F is true if and only if the object functionally mapped by $\langle a \rangle$ has the property functionally mapped by $\langle F \rangle$.

A theory like (CC) gives an account of what it is for a belief with a particular content to be true. But it can be seen as an indirect theory of truth for sentence or utterance-token, given, e.g.:

An utterance-token is true if and only if it expresses a true belief.

If so, then a given utterance-token of the sentence-type "snow is white" will be true just when it expresses a true belief to the effect that snow is white.

Do representationalist theories like (CC) or (TC) count as theories of truth? According to the standard introduced in the last chapter, a theory counts as a theory of truth just when it not only incorporates the truisms as part of the theory, but offers an explanation of at least most of those truisms. Whether or not any representationalist theory is ultimately a *successful* theory of truth, it

does count as a theory of the subject by this standard. This is because representationalists can not only incorporate each of historically prominent core truisms into their theory without alteration; they can argue that their theory offers an explanation of the core truisms.

Consider, for example:

Objectivity: My belief that *p* is true if and only if, with respect to the belief that *p*, things are as I believe them to be.

(CC) and (TC) together with their component theories of representation, each offer an explanation of why this truism is itself true—by offering an account of what it is for things to be as I believe them to be. Suppose, for example, that I believe that Oliver is a cat. According to (CC), we can then say that, with respect to that belief, things are as I believe them to be if and only if the object causally mapped by <Oliver> has the property causally mapped by the <cat>. A similar explanation is available for advocates of (TC). In each case, the point is that not only is the representational theory *consistent* with Objectivity, it *explains* that truism—by offering an account of what it is for its right-hand side to obtain. Moreover, the explanations are reductive: each alleges to say what it is for the world to be as I believe it to be.

One might wonder if either theory, however, is even consistent with (instances of) the T-schema, or the principle that

(TS): The proposition that *p* is true if and only if *p*.

After all, (TS) is concerned with the truth of propositions; (CC) and (TC) are concerned with the truth of beliefs. But this problem is illusory. Recall that we can deduce the relevant instance of (TS) from Objectivity. First, we grant what we called Content above: namely, that what is true when I believe that *p* is true is the proposition that *p*. This secures

(B) The belief that *p* is true if and only if the proposition that *p* is true.

Then, plausibly enough we take it that

When I believe that *p*, things are as I believe them to be if and only if *p*.

From which we can deduce (TS) by way of (B), Objectivity and transitivity of the biconditional. Since we have just shown that (CC) and (TS) are consistent with Objectivity, it follows they are consistent with the relevant instances of (TS) as well.

But might not a parallel worry remain? After all, even if (B) can be granted by a representationalist, theories like (CC) clearly must read (B) in the right to left direction. That is, they must take it that the primary bearers of truth and falsity are, as we might put it, *believed propositions*. But that is inconsistent with taking there to be true propositions that are not believed by someone. Many will not be bothered by this consequence. They will agree with Davidson: “Nothing in the world, no object or event, would be true or false if there were not thinking creatures.”¹¹ Davidson’s point is not that human creatures make something true or false; rather, that the items that bear truth or falsity are mind-dependent—they are beliefs or sentences. Such philosophers accept (B) but not (TS). Representationalists who wish to block the consequence, however, can account for the possibility of unbelieving truths by taking a subjunctive reading of (B) as follows:

(UB): The proposition that *p* is true if and only if were the proposition that *p* to be believed, that belief would be true.

Both (CC) and (TC) can be understood as offering at least the beginning of an explanation of the Norm of Belief. Again, the explanations take a reductive form. According to (CC) for example, true beliefs are correct, as we might put it, because they causally map (or their component concepts do) an object as having a property it does have.¹² Of course, this explanation doesn’t explain the type of norm that is in play—instrumental or intrinsic—nor

¹¹ D. Davidson, “The Structure and Content of Truth”, *The Journal of Philosophy*, 87:6 (June 1990), 279–328; the quote is the very first line of the article.

¹² Whether this is a satisfying explanation of the norm is another question. It suffers the same problems any naturalistic reductive explanation of a norm suffers. Of course, a key

does it plausibly capture all that one might want to say about the correctness of a true belief. But that is not necessary in order for the representational theory to count as a theoretical explanation of truth.

Earlier we noted that any property that makes a proposition correct to believe is regulative of any practice aimed at producing belief. Inquiry is one such practice. Accounts like (CC) or (TC), if correct, again offer an explanation: inquiry aims at truth because true beliefs are those that correctly represent the world as it is.

Representational theories of truth therefore, are *prima facie* consistent with our central truisms about truth. They count as a theory of truth. Moreover, their adherents can claim that the theories explain those truisms in that they imply that beliefs which satisfy the truisms do so because they have a particular property.

A commonly cited further virtue of representational accounts like (CC) or (TC) is that they are a component part of an attractively simple way to explain what determines intentional content, and indirectly, the meaning of sentences used to express such content. According to this explanation of content determination, for example, a belief has the content it does in virtue of its truth-conditions. Truth-conditions are the conditions under which a belief is true. Since (TC) gives an account of what it is for a belief to be true, it can be said to offer an account of the general character of truth-conditions. Therefore, given a truth-theoretical account of content determination, (TC) is arguably part of the explanation of what determines the intentional content of belief-states.

From the vantage point of representationalism, therefore, the theory of truth emerges as a helpful explanatory resource for explaining other phenomena of interest, such as intentionality. Some might think this virtue is a vice, in that (they'll insist) representationalism is "really" a theory of intentionality, not a theory of truth. But this objection misses the point of the theory, which is

benefit of (TC) is that one might be able to ground a better explanation of the norm in the concept of a (proper) function.

to provide an account of both. According to representationalism, truth is reductively *explained* in terms of representation, and hence the two concepts are internally related. Otherwise put: *according to a representational theory of truth, Objectivity, (TS) and like truisms don't tell you what truth is.* They are just that—truisms. Only by combining them with deeper principles like REPRESENT and CAUSAL, do we actually give an informative account of its nature. From the standpoint of the representationalist, insisting that, e.g. Objectivity is all we need to explain truth is like insisting to a reductive physicalist that “the mind is the thing which thinks” is all we need to say to explain the mind. The whole point of representationalism is to provide a naturalist *explanation* of truth—in terms of (what else?) representation. And of course, it works the other way around too. Any naturalistic theory of representation—including, obviously, CAUSAL or TELEOLOGICAL—drags in its wake a theory of truth. It does so because any naturalistic theory of representation must account for what it is for beliefs to correctly represent, and to give such an account is part of what it is to give a theory of truth.

None of this is to say that representationalism is the true theory of truth. Nor is it to agree with the dubious claim that the only way to give an informative account of truth's nature is to give a reductive explanation of the truisms. What the above considerations remind us is that representational theories of mind, since they imply that what makes a belief true is the relation it bears to the world, are successors of older correspondence theories of truth.¹³ And second, that far from being a *recherché* topic in metaphysics, the theory of correspondence as the nature of truth is a central component of the contemporary cognitive science research program. This is a program whose boundaries extend well beyond philosophy.¹⁴

¹³ This point is echoed in Frank Jackson's endorsement of a representationalist theory of correspondence in his “Representation, Truth and Realism”, in *The Monist*, 89 (2006), 51–62.

¹⁴ An excellent illustration of the level of detail at which teleosemantic approaches to representation and truth are presently being carried out is D. Ryder's important paper “SINBAD Neurosemantics: A Theory of Mental Representation”, *Mind & Language*, 19 (2004).

Consequently, to think of truth as correspondence is no more or less spooky and no more or less trivial, than to think of the mind as a device for representation.

2. Representationalism and the Scope Problem

There are various objections and challenges one might raise against any particular representational theory of truth, including of course, against the views just presented.¹⁵ But over and above these theory-specific problems, representational theories all face a problem of scope.

Representationalists tend to favor examples involving cats on mats, white snow, red roses, and other middle-sized dry goods. There is a reason for this. Theories like (CC) or (TC) are only plausible wherever we can make the case that our thoughts about Gs are responsive to the antics of the Gs themselves. And when it comes to cats on mats, this case seems easy to make. If my belief that there is a cat on the mat is true, it is a *response* to—what else?—there being a cat on the mat. When things are working as they should, when our cognitive machinery is firing on all cylinders so to speak, human beings are good detectors of cats on mats. This suggests a constraint on representationalist theories. Say that a given mental state has G-ish content if the proposition that is that content has Gs as its subject. A correspondence theory like (CC) will seem likely as theory of truth for such states only when we can establish

Responsiveness: Mental states with G-ish content are causally responsive to an external environment that contains Gs.

In a bumper sticker, *if we are to correspond, we must respond*.

¹⁵ For objections and discussion, see for example, J. Fodor, “A Theory of Content”, in his *A Theory of Content and other Essays* (Cambridge, MA: MIT Press, 1990); Peacocke *A Study of Concepts* (Cambridge, MA: MIT Press, 1992); K. Neander, “Malfunctioning and Misrepresenting”, in *Philosophical Studies*, 79 (1995), 109–41; R. Millikan, “On Swampkinds”, *Mind and Language*, 11 (1996), 70–130.

Moreover, where responsiveness *does* seem plausible, and we have independent reasons for thinking that the content in question is assessable for truth or falsity, it becomes more *likely* that our mental states with G-ish content have that content in virtue of *representing* Gs. Accordingly, it will seem more likely that when I believe such content *correctly*—when cognitively speaking, success has been achieved—what *makes* my belief that, e.g. the ubiquitous cat is on some mat correct is that it accurately represents, in the sense of (CC), an actual cat on an actual mat.

But where responsiveness is not plausible—either because the states in question aren't appropriately causally responsive or because the external environment contains no Gs that can be so causally responsive—then it is less likely that mental-states with G-ish content have that content because they represent Gs. Some other explanation of their content becomes more likely. And—to anticipate the central lesson—if we nonetheless wish to maintain that the relevant mental states are *true*, some other account of what makes them true must be pushed onto the field.

This highlights the fact that representational theories like (CC) are committed to the following two conditions:

- (1) True beliefs map objects that exist and have their properties mind-independently.
- (2) The object and properties that are so mapped are capable of entering into at least indirect causal interaction with our minds.

The first condition is a consequence of the fact that representational views intend their positions to be realist. An object exists (or has some property) mind-independently at some time just when it would continue to exist (or have that property) even if there were no minds that represented it as having that property. Minds themselves, and even artifacts like hats and hammers exist mind-independently in this sense. The second condition is a consequence of the plausible thought, mentioned above, that mental

representation of external objects is partly a causal process, or at least supervenes on such processes.

The scope problem is that it is difficult to see how both conditions can be fulfilled by all the propositions we intuitively believe to be true. If you allow yourself to focus only on examples involving physical objects, it is easy to slip into thinking that what goes for thoughts about cats and cars must also go for everything else. But on reflection, that is highly implausible. Consider propositions like *two and two are four* or *torture is wrong*. Under the assumption that truth is always and everywhere causal correspondence, it is a vexing question how these true thoughts *can* be true. That two and two are four is unimpeachable, but even granting that numbers are objects, how can any thought of mine be in causal contact with something like a number? Numbers, whatever else they turn out to be, are presumably not objects with which we can causally interact. Moral propositions represent a slightly different puzzle: torture is certainly wrong, but it is difficult to know how wrongness—even if we grant that it is a property—can be a natural property with which we can causally interact.

It is not just condition (2) that is the problem. Even if we retreat from naturalistic theories of representation, condition (1) is sufficient to make trouble all on its own. This is because even non-naturalist correspondence theorists take the facts to which true beliefs correspond to be mind-independent in character.¹⁶ And that alone causes a scope problem. Consider legal propositions like bribing an official is illegal or that flag-burning is constitutionally protected, or that torturing prisoners violates the Geneva Conventions. While we certainly think *our* lawyer should tell us

¹⁶ Examples of such versions of correspondence include A. Newman's *The Correspondence Theory of Truth* (Cambridge: Cambridge University Press, 2002); Vision's *Veritas* and R. Fumerton's *Realism and the Correspondence Theory of Truth* (Chicago, IL: Rowman & Littlefield, 2002). Vision is particularly keen to rebut the charge that suitably robust correspondence theories are hostage to either a particular theory of reference or to the claim that facts or states of affairs must be mind-independent. But this comes at the expense of a theory of truth that involves a particularly permissive (to my mind anyway) theory of reference. See especially 230 ff. See also G. Sher, "In Search of a Substantive Theory of Truth", *The Journal of Philosophy*, 101 (2004), 5–36.

the truth about the law and the Constitution, we don't normally think that in order to be true our legal claims must correspond to some set of entities, the "legal facts"—entities, which, if we are to take such suggestions seriously, must themselves be distinct from the law books, the court rulings and so on. Nor—were there such facts—would it be sensible to see them as mind-independent; for laws themselves are paradigmatic mental constructions. We make laws, and nothing would be legal or illegal if there ceased to be agents whose actions could be assessed as such. How legal facts could therefore be mind-independent seems mysterious.

Of course, advocates of representational theories of truth like we've been discussing are well aware of these examples. Indeed, the history of twentieth-century philosophy is replete with isms that have been posed to deal with them—from attempts to "locate" troublesome entities like numbers in the physical world to theories which instead "eliminate" such entities from their ontologies while trying to preserve, in some way or other, our right to talk as if they existed.¹⁷ But far from constituting an answer to the scope problem, the perceived need for expressivism, fictionalism, error theory and the like is an *acknowledgement of the seriousness of the scope problem facing representational theories*. That is, to claim, as some expressivists have done, that ethical judgments aren't capable of being true or false because they fail to represent objective states of affairs, is just to acknowledge that representational theories of truth fail to be plausible in some domains.

A representationalist can try to avoid the scope problem by watering down his theory of course: "a proposition corresponds to reality just when things are as that proposition says they are" would be one example. But that just brings us back to where we started: a vacuous platitude that any theory of truth can accept. The more substantive the correspondence theory becomes—as when it is seen as part of a larger theory of representation—the more it

¹⁷ The vocabulary of location and elimination is due to F. Jackson, *From Metaphysics to Ethics: A Defense of Conceptual Analysis* (Oxford: Oxford University Press, 1998).

is vulnerable to the scope problem, and the less plausible it is as a universal theory of truth.

3. Superwarrant and Antirepresentationalism

I noted above that correspondence theories of truth have typically been described as “realist”. Likewise, their opponents—such as the pragmatist’s theory of truth—are labeled as “antirealist”. But the last section suggests that a more telling contrast is between those theories that (a) take truth to be a matter of correctly representing a mind-independent world of objects, and (b) those that define truth independently of representation and related notions like denotation and reference. Having just reviewed the prospects of views of the first sort, I now turn to views of the second.

Where representational theories privilege the Objectivity truism, antirepresentational theories typically give pride of place to the idea that truth is the End of Inquiry. Indeed, one of the most well-known versions of the theory, Peirce’s pragmatist view of truth, simply identifies truth with that end: “The opinion which is fated to be ultimately agreed to by all who investigate is what we mean by truth”.¹⁸ Thus, as it is often glossed:

(P) The proposition that *p* is true if and only if the proposition that *p* is fated to be accepted at the end of inquiry.

Peirce’s definition flips the realist’s idea on its head. Rather than saying that we agree on what is true because it is true, Peirce’s thought is that what is true is so because we agree on it. No mention is made of our thought’s having to represent or correspond to some independent world of objects. There may be such a world, but if so, truth is shorn free of it on this account.

Peirce’s position has been subject to numerous revisions and alterations over the past few decades. Thus Hilary Putnam’s “internal realist” defines truth as

¹⁸ C. Peirce, “How to Make our Ideas Clear”, *Popular Science Monthly*, 12 (1878), 286–302.

(PT): The proposition that *p* is true if and only if the proposition that *p* would be warranted to believe in ideal epistemic circumstances for assessing the proposition that *p*.

This is a significant improvement over (P) in at least two ways. First, the use of the subjunctive makes it clear that the view is not in fact committed to the actuality of ideal epistemic circumstances; and second, it takes such circumstances to be not global but tailor-made for each individual belief.¹⁹ Instead of talking about ideal epistemic conditions for all beliefs we say that a belief is true when it would be justified in ideal epistemic circumstances for assessing that particular belief. For my belief that my cat is on the porch, those circumstances would include my being able to see the cat, the light being good, my vision being 20/20 and so on.

But Putnam's view notoriously faces its own problems. One such problem is that the view founders on the so-called conditional fallacy.²⁰ This is a problem that can plague attempts to define a categorical statement in terms of a subjunctive conditional. How the problem applies here can be brought out if we take the proposition in question in (PT) to be

(not-I): Ideal epistemic circumstances for assessing this proposition will never obtain.

Substituting (not-I) for “*p*” in (PT), we arrive at the conclusion that (not-I) is true if and only if it would be warranted in ideal circumstances for assessing (not-I). But if it were warranted in such circumstances it would be false. So, intuitively, (not-I) can only be true if it is false, if (PT) is our theory of truth.

¹⁹ H. Putnam, *Realism with a Human Face* (Cambridge, MA: Harvard University Press, 1990), p. vii; for the original statement of the view, see H. Putnam, *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981).

²⁰ The objection was first posed by A. Plantinga, “How To Be an Anti-Realist”, *Proceedings and Addresses of the American Philosophical Association*, 56 (1982), 47–50. C. Wright generalizes it; see his “Minimalism, Pragmatism, Deflationism, Pluralism”, in M. P. Lynch (ed.), *The Nature of Truth*, 767; Wright, *Saving the Differences* (Cambridge MA: Harvard University Press, 2003), 120–2.

How might such problems be avoided? One reasonable idea, developed by Crispin Wright, is to define a true proposition not as one that would be warranted to believe in ideal conditions, but as one that is warranted to believe in the ordinary sense and remains warranted no matter how our information is expanded or improved. Wright calls this notion “superassertibility”, and applies it (as the name indicates) to assertions.²¹ Rather than discussing his formulation of that idea here, it will suit our purposes to instead use a related notion, which I will call *superwarrant*. Thus, we can take our antirepresentationalist as claiming that

(SW): A belief is true if and only if it is superwarranted.

And we can define superwarrant as follows:

Superwarrant: The belief that *p* is superwarranted if and only if the belief that *p* is warranted without defeat at some stage of inquiry and would remain so at every successive stage of inquiry.

Again, the idea is that a superwarranted belief is one that continues to be warranted throughout an indefinitely long investigation—throughout every successive stage of inquiry. A brief explanation of terms: A *stage of inquiry*, as the name suggests, is a state of warranted information or evidence available in principle in the actual world to some open-minded, receptive inquirer. And a *belief is warranted without defeat at a stage of inquiry* as long as any defeater for the belief at a given stage is itself undermined by evidence available at a later stage. In a sentence: *to be superwarranted is to be continually warranted without defeat*.²²

So superwarrant does not posit an idealized “End of Inquiry”. A superwarranted belief is one that is warranted by some state of

²¹ See his *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992); C. Misak, *Truth and the End of Inquiry* (Oxford: Oxford University Press, 2004) develops a view she identifies as Peircian that has some interesting similarities to Wright’s view and the account of superwarrant below.

²² Compare Beall and Restall’s account of constructivism in their *Logical Pluralism* (Oxford: Oxford University Press, 2005).

information available to *ordinary inquirers*, which, in fact, would never be defeated or undermined by subsequent increases of information also available to ordinary inquirers. Moreover, superwarrant is a stable property: if a belief is superwarranted, then it is superwarranted at any stage of inquiry.

Beyond this point, sympathizers are apt to disagree over what else to say. One source of disagreement will be over whether to take inquiry as strongly incomplete: that is, to hold that successive eternal inquiry will neither warrant a belief nor its negation. Coupled with (SW) this entails that we do not accept the law of excluded middle or the thought that

P or not P.²³

Another source of potential disagreement is over the nature of warrant. Since (SW) constructs truth out of warrant, the theory is simply incomplete unless we are told what it means for a belief to be warranted. Presumably, antirepresentationalists won't be attracted to a traditional foundationalist approach, nor to an externalism which takes warrant to result from our reliable responses to an external world of objects. As we will have occasion to see in greater detail in Chapter 8, a more natural fit is the coherence theory of warrant. According to this view, roughly, a belief is warranted just when that belief is a member of a system of beliefs which exemplifies what we might call the coherence-making features: mutual relations of deductive and inductive support, simplicity, predictive power and consistency. Considerably more must be said about the theory of course; but for present purposes, I will trust that this intuitive characterization will serve. In any event, if this is how warrant is understood, (SW) will amount to a coherence theory of truth. Such a theory might look like this for example:

(SC): The belief that p is true if and only if that belief is supercoherent.

²³ Together with associated inferences, such as double-negation elimination.

If we see (SW) as a recipe of sorts, we then might define super-coherence as:

A belief is supercoherent just when it is a member of a coherent system of beliefs at some stage of inquiry which would remain coherent without defeat in every successive stage of inquiry.

This sketch of a super-coherence theory of truth can be seen as one way to begin completing the recipe given by (SW). There may well be others, and I will have significantly more to say about this particular view in Chapter 8. My present point is simply that no theory which defines truth in terms of superwarrant can be complete without an account of what warrant consists in.

Does any version of (SW)—(SC) included—count as a theory of truth? It seems that they do. For along with the representationalist, they both incorporate and claim to explain the truisms. This is most obvious in the case of Norm of Belief and End of Inquiry. For superwarrant, like warrant, is a normative notion. Thus, this allows the antirepresentationalist to argue that true beliefs are correct because they are superwarranted. And since whatever makes a belief correct is regulative of any practice that aims to produce belief, the antirepresentationalist can go on to add that true beliefs are a proper End of Inquiry because they are superwarranted. Indeed, many philosophers have found this last claim particularly plausible, since it has seemed to some difficult to distinguish between the goal of believing what is eternally justified and the goal of believing what is true.²⁴

It may seem that the antirepresentationalist will have a harder time accepting Objectivity. But in fact, the antirepresentationalist can easily incorporate that truism just as long as she accepts two additional platitudes. First, she must accept, as we have throughout, the following truism about belief: that when I believe that *p*, things are as I believe them to be if and only if *p*. And second, she must grant that the metaphysical thesis that we can call *idealism*, namely:

²⁴ See R. Rorty, "Is Truth a Goal of Inquiry?", reprinted in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA: MIT Press, 2001).

(Id): p if and only if the belief that p is superwarranted.

From our truism about belief, together with the metaphysical view of idealism, and with (SW) itself, Objectivity follows. And it hardly seems that the antirepresentationalist can reject idealism in any event, given that it follows from the independently plausible (BS) and (SW).

So antirepresentational theories like (SW) count as theories of truth. Moreover, they have a similar virtue as their representationalist cousins: they can form a component part of an attractively simple truth-conditional explanation of intentional content and the sentences used to express such content. Moreover, and as Dummett, Wright, and others have emphasized, where truth is understood as something epistemic like superwarrant, a truth-conditional account of content is in some ways more plausible.²⁵ For an account of content must presumably be able to explain how we can *grasp* such content and how we can *manifest* that grasp in our behavior. On a truth-conditional account, to grasp a proposition is to understand the conditions under which that proposition is true. And where truth is superwarrant, this involves grasping the conditions under which belief in that proposition would be warranted. And this has seemed to some to be a more manageable feat. But whether or not one agrees, (SW) at least offers to play a crucial role in a theory of content.²⁶

4. Antirepresentationalism and the Scope Problem

The fact that (SW) or (SC) count as a theory of truth and have the virtues of a substantive theory—it allows us to use truth as an

²⁵ See M. Dummett, *The Seas of Language* (Oxford: Oxford University Press, 2003); C. Wright, *Realism, Meaning and Truth* (Oxford: Blackwell, 2001).

²⁶ As I'll note below, even if one doesn't think this explanation of content is attractive across the board, it can seem like an especially attractive explanation of the content of beliefs about matters where it is implausible that we are in causal contact with entities that make our beliefs about them true.

explanatory resource for other phenomena—does not mean that it is the correct theory of truth. Like its representationalist rival, it faces a problem of scope.

Not surprisingly the problem arises for quite different reasons. First, and most obviously, (SW) requires that *all* content is non-representational. Earlier, we noted that representationalism is viable only where Responsiveness is satisfied. But we also noted that where it *is* satisfied, representationalism is highly plausible. That is, where (a) mental states whose content is G-ish are responsive to an external environment that contains physical Gs, and (b) we have independent reason for thinking that content to have truth-value, those mental states are likely to have their content because they represent Gs. (Again, we can compare perception: a strong reason for taking our perceptual states to be representational is that such states are caused by an external environment, and we have independent reasons for taking such states to be capable of being veridical or not.) And surely there are some kinds of beliefs for which Responsiveness *is* satisfied. Beliefs about the middle-sized objects with which we are in perceptual contact seem like prime candidates. For such beliefs, it is much more plausible to think that they are true or false in virtue of whether they correctly represent what they are about. And hence it is implausible that their truth will consist in being superwarranted.

Second, (SW) implies that truth is globally epistemically constrained.²⁷ And this seems implausible. The implausibility in question can be neatly brought out by a modification of reasoning familiar from Fitch's knowability paradox.²⁸ Take a proposition like

(F): P and no one will ever have any warrant for P.

Given (SW) and the T-schema, one can derive

²⁷ This is a point repeatedly emphasized by Wright (although often it seems on deaf ears); see his *Truth and Objectivity* and "Minimalism, Deflationism, Pragmatism, Pluralism".

²⁸ The argument here is inspired by T. Williamson, "A Critical Study of Truth and Objectivity", *International Journal of Philosophical Studies*, 30 (1994), 130–44.

(F) is superwarranted if and only if P and no one will ever have any warrant for P.

Now clearly (F) cannot be superwarranted.²⁹ Hence from that fact and the above

It is not the case that (F)—that is, it is not the case that: P and no one will ever have any warrant for P.

Which is to say that, if P, then someone will at some point—at some stage of inquiry—have warrant for P.

If we were to focus on only some kinds of truths, this consequence may not seem too bad. Consider, to use Crispin Wright's favorite example, truths about what is or isn't funny. It is odd to think that a joke is funny even if no one will ever have warrant in believing that it is—even if, nobody ever laughs, in other words. Likewise for legal truths. It is difficult to see how a proposition of law might be true even if no evidence is ever available for (or against) it—even in principle. For otherwise, it would be possible for there to be a true proposition of the form "x is illegal" even if no one would ever be warranted in believing that it is, or is not legal. And that in turn means that there could be unknowably illegal actions—actions I might even be doing right now. But that seems absurd. As we'll see in Chapter 8, a similar point is sometimes made in the case of moral truth.

So an epistemic constraint seems plausible in the case of some normative truths at least. Be that as it may, it seems highly implausible that such an epistemic constraint could be motivated across the board. Surely, one might think, there are at least *some* truths—perhaps about the distant past, or far side of the universe, or the number of stars right now, for which no evidence will ever be available in principle. Humility in the face of the size of the universe seems to demand that. And yet (SW) would

²⁹ Argument: (F) can only be superwarranted if both of its conjuncts are. But both conjuncts cannot be superwarranted. For if the first is superwarranted, then the second conjunct is obviously false, and by (SW), not superwarranted.

seem to require us to deny this. And this seems implausible: surely there can be *some* truths for which we will never have any warrant.

The problem I want to highlight here can again be put in terms of a dilemma: Either the antirepresentationalist admits that her theory has an absurd consequence—in this case, the consequence that all truths are at some point justifiably believed by someone—or admit that her view has limited scope. If she does the latter, she can hold that the truth of some propositions does transcend any evidence we could have for them, and such propositions will include of course propositions like (F).

The above argument is not the only way to force a dilemma of this sort onto the antirepresentationalist. For what we might call the idealist character of antirepresentationalism also leads to a parallel choice between an absurdity and admission that the view must be limited in scope. And here the ensuing dilemma spells a problem even for a more specific version of (SW) like (SC)—which otherwise might be thought to elude our modified Fitch objection.

The oldest objection to coherence theories, first formulated by Russell, is the so-called “many systems objection”. The thought is that there could be more than one supercoherent system of beliefs, and that, e.g. P and $\sim P$ could be members of such rival systems.

One advantage of our Wrightian approach—to (SC) in other words—is that it is not immediately clear that this objection applies: for a belief is supercoherent only if it is a member of a system that remains coherent in the face of all potential increases in information—including information from rival systems of beliefs. So it is unclear whether there could be more than one such theory in advance of hearing more about supercoherence and its nature. In any event, let us assume for the sake of argument—and only for the moment—that *the many systems objection can be answered*. Even if we do so, further reflection shows how, if (SC) is taken as a global theory of truth, the theory must either embrace an absurdity

or acknowledge that some beliefs are true in some other manner than being supercoherent.³⁰

So, assume that there is at best one possible supercoherent system—where by this I simply mean that system that remains continually coherent no matter how enlarged and so on. Allow our coherentist to pick out whatever system S he thinks is the best candidate. We can now abbreviate our theory as:

(SC*): The belief that p is true if and only if the belief that p is a member of S.

Now consider the claim that

(I): S is the one and only supercoherent system.

Now according to (SC*), what would make (I) true if it were true? That is, what makes S—that very system—the one and only supercoherent system *if it is*? According to our theory, all truths are so precisely in virtue of their being members of supercoherent systems. And by hypothesis, there is one and only one such system. So if (I) is true, its truth must consist in its being a member of the one and only supercoherent system of beliefs. And (I) claims that S is that system. So if (I) and (SC*) were both true, then (I) would have to be a member of S. We can put this by saying that if (I) is true, then S must say of itself that it is the one and only supercoherent system of beliefs. And crucially, the converse also seems to hold: if S says of itself that it is the one and only supercoherent system of beliefs, then this must be so. And this appears to be an embarrassing consequence. For it is absurd to think that any system of beliefs could be supercoherent *just because* it says of itself that is so. Believing my belief system is supercoherent doesn't make it so.

We can spell this out as follows. As we noted above, a theory like (SC*) can only incorporate Objectivity into their account by adopting a principle like (Id) above, that is

³⁰ For a different development of a similar argument as the one that follows, see R. Walker, *The Coherence Theory of Truth* (London: Routledge, 1989).

p if and only if the belief that p is supercoherent.

Or, as we might now put it:

(Id*): p if and only if the belief that p is a member of S.

Of course, one might point out that any such equivalence can be “read” as it is sometimes put, in one of two ways. The right to left, or Socratic reading of the biconditional would have it, in effect, that the belief that, e.g. roses are red is a member of S in virtue of roses’ being red. But that is clearly not the intended reading. For if we were to read (Id) in that fashion, we would have to read (SC*) that way too. And that would violate the spirit of the account, which is to explain truth in terms of supercoherence, not the other way around. Hence, the advocate of (SC*) will read both biconditionals left to right, or *Euthyphronically*. That is, she will hold that

(SCV): The belief that p is true in virtue of the belief that p being a member of S.

And

(IdV): p in virtue of the belief that p being a member of S.

And this would presumably not engender any complaints from actual advocates of the coherence theory of truth, many of whom hardly shirked the idealist label in any event.³¹

It should be clear that the embarrassing consequence follows quickly. This can be seen by simply substituting (I) in for “p” in (IdV). That is,

(IdV¹) S is the one and only supercoherent system in virtue of the belief that S is the one and only supercoherent system being a member of S.

³¹ See for example, B. Blanshard, *The Nature of Thought*, vol. 2 (New York: Harper Collins, 1939), 260–79; F. H. Bradley, *Essays on Truth and Reality* (Oxford: Clarendon Press, 1914). See also R. Walker’s helpful study of the coherence theory for discussion of this point, “The Coherence Theory”, in Lynch (ed.), *The Nature of Truth*, 123–58.

Which is just our embarrassing consequence: if S says of itself that it is supercoherent, then it is.

Some readers may wish to object that I have mischaracterized the coherentist position. It will be insisted that if S is to be the supercoherent system, then, *in addition* to S saying of itself that it is supercoherent, S *must really be supercoherent*. In other words, the thought goes, the correct principle is not (Id*) but

(Id**): p if and only if the belief that p is a member of S and S is the one and only supercoherent system.

Likewise for (SCV) and (IdV): we should understand each as employing (I) as part of a conjunction on their right-hand side.

But it is difficult to see how this move avoids the essential problem. The relevant conjunction will be true only if both conjuncts are true. And for each conjunct we can ask what makes it true if it is. In particular, we can ask what makes it true that S is the one and only one supercoherent system. And that is just our original question: if (I) is true, what makes it so?

The advocate of (SC*) would seem to have only two choices. First, she can give the same answer again: namely, she can appeal to her theory of truth: if (I) is true, what *makes* (I) true (whether or not it is part of a conjunction like the above) is that it is a member of S. And given that what makes p true is sufficient for p, this means it is sufficient for S being the supercoherent system that it say of itself that it is, which is just the embarrassing consequence. Moreover, once this route is taken, it is difficult to see how the coherentist can—contrary our initial assumption—solve the “many systems” objection. For according to the reasoning just given, there is nothing but coherence all the way down. Ultimately, nothing determines that a system is supercoherent other than its saying of itself that it is. So how can the coherentist disallow the fact that any system that says of itself that it is supercoherent *is* supercoherent? Only it seems, by insisting that there is some fact that consists in

something other than supercoherence, a fact that determines what is supercoherent.³²

Thus the coherentist's other choice: she can easily avoid our argument by simply accepting that propositions like (I) and the like are not true in virtue of being members of S. But of course to do so is to abandon her theory and admit that not all truths are explained in terms of supercoherence. This latter point was anticipated by Russell:

... the objection to the coherence theory lies in this, that it presupposes a more usual meaning of truth and falsehood in constructing its coherent whole, and that more usual meaning, though indispensable to the theory, cannot be explained in terms of the theory.³³

Precisely, one might think, what has happened here: the scope problem has again reared its head, and this time with a vengeance. (SC), on pain of absurdity, simply cannot be a theory of truth for all propositions. Indeed, a moment's reflection suggests that the argument generalizes. For no matter what account of warrant she gives, the antirepresentationalist must say something about what makes it the case that a belief is superwarranted, and she must avoid the answer that it is just superwarrant "all the way down".

³² Might it help to insist that in order to be ideally coherent, S must not merely say of itself that it is supercoherent—it must meet further constraints? No. For consider the belief

(A) System S meets constraint C (has beliefs that are mutually explanatory or consistent or whatever).

Here the problem simply repeats itself. What makes (A) true? It cannot simply be a fact not explained in terms of coherence. Accordingly, if (I) is true, and S is the supercoherent system, then (A)'s truth must consist in its being a member of S. And if (I) is not true, then by the terms of our suggestion, neither is (A). Note that this conclusion is not altered by any particular fleshing out of what meeting constraint C amounts to. And nor would it help to insist that membership in S is not itself sufficient for S to meet the relevant constraints. For whatever further constraints might be imagined to be required for supercoherence, so long as (SC) is taken seriously, we again face only two options. We can either say that what makes (A) true is something other than coherence, or that (A) is true just because it is a member of S, that is, we can admit that according to (SC) it can be true that those constraints are met if and only if S says of itself that it has met them. Which—given that meeting those constraints is sufficient for being supercoherent—in turn means the belief that S is supercoherent is true if and only if S says of itself that it is supercoherent.

³³ "The Monistic Theory of Truth", reprinted in his *Philosophical Essays* (London: George Allen & Unwin, 1966), 136.

Taken together, the above arguments show that the scope problem afflicts antirepresentational theories just as it does representational theories of truth. Moreover, the force of the problem is arguably greater in this case. For the arguments just given show that lest they embrace an absurdity, antirealists must hold that there are some propositions for which (SW) cannot be the correct theory of truth.

5. A New Strategy

Traditional theories of truth are attractive precisely because they aren't simply a collection of truisms. They attempt to explain our preconceptions about truth—by giving a single unifying metaphysical theory of truth's nature. In so doing, they hold out the possibility that we can use this theory to understand other philosophically interesting phenomena—such as intentionality and meaning—which seem closely linked to truth. But if the arguments given above are on the right track, the two most venerable substantive theories of truth face a grave problem of scope. And the history of the debate over truth makes it likely that this problem is in fact entirely general: for any sufficiently robustly characterized truth property *F*, there appears to be some kind of propositions *K* which lack *F* but which are intuitively true (or capable of being true).

As I noted in the Introduction, philosophers typically adopt one of two strategies in response to this problem. Those adopting the first strategy hold fast to their favored theory of truth and damn the counterexamples. They simply deny that various troublesome propositions are true (error theory), or even capable of being true (expressivism). The second strategy dismisses the whole project of giving a metaphysical theory of truth, and declares that all propositions are equally apt for truth in an entirely thin sense. On this approach we give up on trying to explain other interesting phenomena in terms of truth and the concept of truth is taken only as an expressive device. This is the deflationary strategy.

In the rest of this book I aim to offer a third strategy. At heart, it consists in this thought: that there can be more than one property that can make beliefs true. Neither correspondence nor superwarrant can be the nature of truth *simpliciter*. But each holds out a possible way in which certain kinds of beliefs and their contents can be made true, providing, as we'll see, those beliefs and contents meet certain conditions. If this broadly pluralist approach to truth could be made coherent, then there is more to say about truth than the deflationist believes, but the more there is to say depends on the type of proposition in question. We would retain truth as an explanatory resource for other phenomena—including, for example, intentional content. Moreover, it would then be possible to explain the *diversity* of intentional content. The explanation would consist in giving distinct explanations of the conditions under which our belief's contents are true. Moral beliefs, for example, might be true, and therefore have the content they have, in virtue of being superwarranted, while beliefs about physical objects might be true by corresponding to the facts about those objects.

The hope, in short, is that we might heed Wittgenstein's command to mind the differences between forms of thought and yet still hold onto the idea that we can have true beliefs about morality, or economics, or mathematics. We would have content diversity yet cognitive unity.

3

Truth as Many

There are many kinds of eyes—even the Sphinx has eyes.
And there are many kinds of “truths” and consequently there
is no truth.

Nietzsche, *Will to Power*

1. Local Truth

Traditional theories of truth assume that truth has a real essence—a substantive property common to all truths, no matter what the subject. One lesson of the last chapter is that it is exceedingly difficult to find that common property in the face of the sheer diversity of our thought. Theories that seem plausible when applied to beliefs about the physical world around us (such as the correspondence theory) are less plausible when applied to beliefs about norms. And theories that seem plausible when applied to the beliefs about norms (such as the coherence theory) seem much less convincing when applied to beliefs about the physical world.

Seen in this light, it is natural to wonder whether the theories only apply *locally*. Perhaps a representationalist theory is correct for some types of belief contents, while the coherence theory is correct for other types. Considerations raised in the last two chapters, while far from exhaustive, suggest the general shape of this idea. A theory counts as a theory of truth, we’ve said, when it explains the nominal essence of truth—that is, most of the truisms about truth. Both the representationalist theory and the superwarrant theory

of truth count as truth-theories. But it is one thing to count, another to be viable. One fall-out of the scope problem is that these theories are viable only where certain additional constraints are met. Thus, for example, I argued that representationalism is plausible only when we can make the case for the responsiveness of our mental states—or the idea that mental states with G-ish content are causally responsive to an external environment that contains Gs.

In a similar vein, Crispin Wright has claimed that a domain of thought (he would say “discourse”) is representational just when it exhibits what he calls *cognitive command*. Domains of thought do so when it is a priori that differences of opinion formulated within that domain, saving those that can be excused because of vagueness, “involve something which may properly be regarded as a cognitive shortcoming”.¹ In other words, if it makes sense to say that truth in some domain is a matter of correct representation of facts and properties in the world, then disagreements must stem from one party or the other (or both) *misrepresenting* those facts and properties. Someone has not paid attention to the evidence, or misperceived the facts, or calculated the probabilities incorrectly.

These constraints serve to underline a familiar thought: that mental states are representational only where those states display a particular *direction of fit*: one that goes from world to mind. That is, representational states are those that are responsive to an objective independent world of objects, a world that they can fit or fail to fit.

Our discussion of superwarrant, on the other hand, suggests that it is viable as a theory of truth for a domain only if it meets two very different conditions. First, the concepts that we employ within the domain must impose an *epistemic constraint* on the truths of the domain. That is, where superwarrant is

¹ C. Wright, *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992), 144; see also his “Truth in Ethics”, in *Saving the Differences* (Cambridge, MA: Harvard University Press, 2003), 198.

taken for truth, it must be in principle possible for someone at some time to have warrant for believing any given proposition. Unless some plausible independent argument can be given for this supposition, truth can't be superwarrant.² Second, the domain must be *non-representational* in character. That is, the propositional content that composes the domain must be the object of non-representational mental states. We noted that representationalism is viable only where responsiveness is satisfied. But we also noted that where it *is* satisfied, and we have independent reasons for thinking that the relevant contents have truth-values, it is likely that representationalism will be correct. That is, where mental states whose content is G-ish are responsive to an external environment that contains physical Gs, those mental states are likely to have their content because they represent Gs. And where there is representation, there is correct and incorrect representation, which is to say that truth is better thought of in terms of correspondence, not superwarrant.³

In the last chapter, I argued that these additional requirements will not be met across the board for either representationalism or superwarrant. If we take our everyday ascriptions of truth seriously, we must admit that globally speaking, truth isn't superwarrant or representation. Nonetheless, it may be that the requirements can be met locally. If we believe that they can, we are committed to what I'll call pluralism about truth. So far we've only described this view in the roughest of terms. We've talked variously of different "theories of truth" and the idea that there is more than one way for propositions to be true. It is time to get clear as to what this means.

² This is a point repeatedly emphasized by Wright (although often it seems on deaf ears); see his *Truth and Objectivity* and "Minimalism, Deflationism, Pragmatism, Pluralism".

³ Many philosophers hold that beliefs are *by definition* representational. To believe just is to represent the world to be a certain way. If so, and if, as I just urged, superwarrant is viable as a local "theory of truth" only where representationalism is not, then we must conclude either that beliefs are not what is true or false in domains where superwarrant is a plausible theory of truth, or that there is more than one way to be a belief—that is, some beliefs are not representational. I will suggest that the proper route is the latter in Chapter 7.

2. Simple Alethic Pluralism

Here's one thing it shouldn't mean: that "true" is simply ambiguous in the way "step" or "bank" is ambiguous. That is, the word conveys different concepts or meanings in differing contexts. Sometimes "truth" means "correspondence" (in the teleofunctional sense, say); sometimes it means "superwarrant". Call such a view simple alethic pluralism (or SAP). I'm not sure anyone actually advocates SAP, but a glance at the literature, together with anecdotal evidence, suggests that lots of folks seem to think alethic pluralists must be committed to it.⁴ They are wrong; but given the widespread belief that things are otherwise, it will help to get clear on what's wrong with SAP, whether or not anyone professes to hold it.

The reasons for rejecting SAP range from something of a cheap-shot to the level of profound problem. The cheap-shot, to quote Kripke, is that "it is very much the lazy man's approach in philosophy to posit ambiguities when in trouble"⁵ This is not to say that positing ambiguity isn't ever helpful in philosophy. But it is justified only on the basis of serious theoretical pressure. And there isn't any such pressure here. Indeed, what pressure there is goes in quite the opposite direction, as the following substantive problems indicate.

Before getting to these problems, some clarifications: SAP is the idea that there is more than one concept of truth, that truth-talk is equivocal, and therefore that the meaning of "true" is, in a sense, context-sensitive. This needn't imply that truth is a relative

⁴ M. Sainsbury, "Crispin Wright: Truth and Objectivity", *Philosophy and Phenomenological Research*, 56 (1996), 899–904; P. Pettit, "Realism and Truth: A Comment on Crispin Wright's Truth and Objectivity", *Philosophy and Phenomenological Research*, 56 (1996), 883–90; originally interpreted Wright as embracing SAP; for recent allegations of the same kind see W. Kühne, *Conceptions of Truth* (Oxford: Oxford University Press, 2003).

⁵ S. Kripke, "Speaker's Reference and Semantic Reference", reprinted in A. P. Martinich (ed.), *The Philosophy of Language* (Oxford: Oxford University Press, 1990), 19; original appeared in *Midwest Studies in Philosophy* 2 (1977).

predicate of course, or that “true” means something like “true for”, although it doesn’t rule that out either. But if we assume, as one might expect, that the various concepts in question each pick out a different property, it does imply that there is more than one property of truth.⁶

To this we might add the following further complication. Views about the meaning of “true” bring in their wake views about the meaning of “false”. But a little reflection shows that the simple alethic pluralist has two options about falsity. First, she may say that there is a single way of being false: namely, a proposition is false when it is not true in any sense (or that is neither true₁, nor true₂ nor true₃ etc.). This position, however, threatens to make it rather more difficult for a proposition to be false than we ordinarily think, since it implies that a proposition is false only if it lacks every property expressed by the truth-predicate. Alternatively, she might claim that there are as many concepts expressed by “false” as there are concepts expressed by “true”, for it is plausible that to say that “x is false” in some domain implies “x is not true” in that domain—that is, that x lacks the property (whatever it is) ascribed by the word “true” in that domain. Unless I say otherwise, I shall take this latter position to be that of the simple alethic pluralist.

The first problem for SAP is how to understand *mixed inferences and validity*.⁷ According to a standard way of understanding validity,

⁶ As we’ll see, the reverse doesn’t necessarily hold. One might think that “true” denotes different properties in different domains without thinking that there is any change in the concept.

⁷ There is a growing literature on this problem. C. Tappolet, “Mixed Inferences: A Problem for Pluralism about Truth Predicates”, *Analysis*, 57 (1997), 209–10; Sainsbury, “Crispin Wright: Truth and Objectivity”, 899–904; M. P. Lynch, *Truth in Context* (Cambridge, MA: MIT Press, 1998), ch. 5. See also “A Functional Theory of Truth”, in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA: MIT Press, 2001), 723–50; and J. Dodd, “Recent Work on Truth”, *Philosophical Books*, 43 (2002), 279–91; N. Pedersen, “What Can the Problem of Mixed Inferences Teach us about Alethic Pluralism?”, *The Monist*, 89:1 (2006), 3–117; D. Edwards, “How to Solve the Problem of Mixed Conjunctions,” *Analysis*, 68.2 (2008), 143–9; A. Cotnoir, “Generic Truth and Mixed Conjunctions: Some Alternatives”, *Analysis*, 69:2 (2009).

validity preserves truth. That is, a valid argument is one where if the premises are true, the conclusion must be true. Now consider the argument that

If you hold a prisoner indefinitely and without charge, you violate his rights. This prisoner has been held indefinitely and without charge. Therefore, this prisoner's rights have been violated.

The second premise of this argument is a claim about the physical facts of a prisoner's incarceration. The conclusion is a normative claim. So if it means something different to say that the second premise is true than to say that the conclusion is true, there is no single property being preserved from premises to conclusion in this argument. As a result, the advocate of SAP must either explain validity in some less than standard way, or she must admit that there is a univocal concept of truth after all.⁸

A related problem concerns the truth of compound propositions. Consider the proposition that

Murder is wrong and two and two make four.

Intuitively, the conjuncts of this proposition are from very different domains. What explains, then, the truth of the conjunction itself? In response, the advocate of SAP may say: A conjunction is true just when its conjuncts are both true in some sense or other. Perhaps, but this reply begs the real question, which concerns not the conjuncts but rather the sense in which the conjunction *itself*

⁸ Jc Beall has argued that the pluralist can get around this problem by appealing to a many-valued logic; see his "On Mixed Inferences and Pluralism About Truth Predicates", *The Philosophical Quarterly*, 50 (2000), 380–2. According to Beall, the logic needn't give up on the standard account of validity as necessary truth preservation because it can appeal to a concept of designated value, where every way of being true is a designated value. Thus an argument is valid on this approach just when if all its premises are designated, so is its conclusion. But why should we care whether an argument is "designation-preserving"? One is tempted to say: because such arguments preserve *truth*. In any event, note that this solution does not help with the problem of mixed compounds; see Tappolet's 2000 reply to Beall, "Truth, Pluralism and Many-Valued Logics", *Philosophical Quarterly*, 50 (2000), 382–5.

is true. This is a problem not just for conjunctions but for the truth of disjunctions and conditionals as well.⁹

Generalizations involving truth pose a third set of problems. A useful fact about the concept of truth is that it allows us to make blind generalizations. That is, we can say: Everything Socrates said was true. But Socrates said lots of different sorts of things. So if “true” means different things when predicated of different sorts of propositions, we face two related questions about such generalizations. First, in what sense is the generalization *itself* true? Since it might range over all sorts of propositions of various domains, it would be arbitrary to say that it is true in any one of the senses that a particular proposition that it ranges over might be true. And surely it isn’t true in some special “general” sense. Second, how do we even understand such generalizations? Perhaps it means: everything Socrates said was true in some sense or other. That is, that for every proposition Socrates said, that proposition is either true₁ or true₂ or true₃ and so on. But this seems unsatisfying. For one thing, it simply seems to miss the original point of the generalization, which doesn’t say anything about a disjunctive property of Socrates’ actual and possible claims. And if the simple alethic pluralist takes it that there is more than one way for a proposition to be false, then she would likely have to concede that much of what Socrates said was false in some sense or other as well. And *that* definitely seems to miss the point of the generalization.

A fourth problem concerns whether SAP can account for the normative truisms we canvassed in Chapter one: what we called Norm of Belief and End of Inquiry. Both norms are general; and this generality poses a particular problem for SAP with regard to

⁹ Williamson first raised this problem against Wright, “A Critical Study of Truth and Objectivity”, *International Journal of Philosophy*, 30 (1994), 130–44; See also Tappolet (2000) who calls this the problem of mixed conjunctions; but it seems apparent that the problem extends farther than conjunctions. It might be thought that the truth of a mixed disjunction could be handled by saying that such a disjunct is true just when at least one of its disjuncts is true in some sense or other. But again, we might ask: in what sense is the disjunction itself true? Similarly for conditionals. Negations, presumably, can be handled by saying that the negation of a proposition is true or false in the same sense that the original proposition is true or false.

taking truth as a norm of belief. As we noted in Chapter 1, the normative generality here can be illustrated by Dummett's well-known analogy between the concept of truth and the concept of winning a game. You wouldn't understand a notion of a competitive game if you were not aware that the point of the game is to win. That winning is the point of competitive games makes them the type of games they are. As such, winning is a univocal, single norm of any competitive game. Analogously, that the standard of correctness for beliefs is truth—that truth is the "aim of belief" is part of what makes beliefs what they are; truth is a constitutive norm of *any* doxastic practice. Yet this would be impossible if there were no single concept picked out by the word "true". If there were no single concept capable of being expressed by "true" in every domain, there would only be a list of particular norms, e.g. in the moral domain, one should believe what is true₁; in the mathematical domain one should believe what is true₂. This overlooks the fact that it is part of the point of the truth concept that it applies to all types of beliefs because it is part of what makes a belief what it is.

It might seem that the advocate of SAP will have an easier time with the idea that truth is an end of inquiry. She can claim that what this truism amounts to on her view is that different forms of inquiry have different ends: in one domain we strive to have beliefs which are true₁, in another, beliefs which are true₂. But here we have an obvious problem: the most obvious explanation for *why* one ought to believe what is, e.g., true₁ is that believing what is true₁ is simply a way of believing what is true. But that explanation means that being true₁ isn't normative itself. It is normative because it is a way of being true.

The simplest point to make about SAP is that it isn't even a pluralist view of *truth* at all. It is a pluralist view of *the meaning of the word* "true". Compare: we don't say that there is more than one way to be a bank. We say there are different meanings to the word "bank". We don't think there is anything in common between riverbanks and the Bank of America except the fact we call them

“bank”. Seen clearly, SAP would seem to fall victim to Nietzsche’s remark that, “there are many kinds of eyes—even the Sphinx has eyes. And there are many kinds of truth, and therefore there is no truth.” In other words, simple pluralism about truth is really a disguised form of truth *nihilism*. If we really took it seriously, we’d just stop talking about what is true and talk about the various properties the word ambiguously picks out.

Taken together, SAP’s problems suggest that any coherent pluralism about truth must stop short of saying that the truth-predicate expresses completely different concepts across domains. If we are to be pluralists, we need a unifying explanation of the different types of truth—an explanation that nonetheless leaves room for truth to be plural in some significant sense, that allows truth to be many but one.

3. Wright’s Reductive Pluralism

We’ve noted that despite its fundamental mysteriousness, there is much about truth that is common knowledge. The truisms attest to this. Might they alone be sufficient to provide the pluralist the needed ties that bind? Crispin Wright, in laying out his own well-known theory of truth, has essentially made just this suggestion. This needs looking at in detail.

Wright’s basic position, which he calls “minimalism”, is the view that an analysis of the concept of truth proceeds with reference to a set of principles (or “platitudes”). These principles include several of our truisms, such as the T-schema and Warrant-Independence, but they also include several plausible additions, such as:

To assert is to present as true.

Truth is absolute and does not admit of degrees.

Truth-apt contents have truth-apt negations and can be embedded within conditionals, conjunctions, etc.¹⁰

¹⁰ For a fuller account of the relevant principles, see C. Wright, “Truth: A Traditional Debate Reviewed”, *Canadian Journal of Philosophy*, 24 (1999), 31–74.

Together, these and similar principles provide what Wright calls an “analytic theory” of the concept of truth. As Wright earlier put it,

The proposal is simply that any predicate that exhibits certain very general features qualifies, just on that account, as a truth predicate. That is quite consistent... with acknowledging that there is a prospect of pluralism—that the more there is to say may well vary from discourse to discourse.¹¹

According to Wright, the specifics about truth in a particular domain hang on the a priori facts about that domain. This is just the suggestion we’ve been entertaining: and indeed, Wright is well known for advocating that in some domains, a proposition is true just when it is *superassertible*, or what I’ve called superwarranted; in other domains, Wright seems to think that something like correspondence would do the trick, although he hasn’t suggested how he understands that concept.

But right here a question arises. For a natural way of reading the above passage takes it to imply that there can be more than one “truth-predicate”. Early commentators took this to commit Wright to more than one concept of truth, and so to SAP. In reply, Wright vigorously protested that his point was not that “true” is ambiguous, and that the whole point of his position was to hold that, “on the contrary, the concept admits of a uniform characterization wherever it is applied—the characterization given by the minimal platitudes”.¹² The form of pluralism relevant to his position, Wright contended, was therefore not SAP, but one that allowed truth to admit of “variable realization” (Ibid). Wright has more recently clarified the point:

Even if the *concept* [of truth] may be fully characterized by reference to certain basic a priori principles concerning it, the question of which *property* or *properties* of propositions, or sentences, realize the concept

¹¹ Wright *Truth and Objectivity*, 38.

¹² Wright “Responses to Commentators”, in his *Saving the Differences*, 101.

can be still be sensibly raised for every discourse in which truth has application....¹³

So: one concept, more than one property. An analogy with definite descriptions and contexts of utterance helps to understand what is going on here. Take a definite description like “the brightest object in the sky”. Depending on the context—time of day, the location and so on, the brightest object in the sky could be the sun, or the moon, or a satellite. Nonetheless, the phrase expresses the same concept on any particular occasion as it does on any other; it isn’t ambiguous in the way “step” or “bank” are. Definite descriptions like “the brightest object in the sky” are univocal, but non-rigid; they pick out different objects in different contexts. Similarly, the description “the color of the sky at noon” univocally picks out different *properties* in different contexts. Intuitively, it differs from a term like “magnetism” which presumably names the same property in every context and, indeed, in every world.¹⁴

A reasonable interpretation of Wright’s view is that it takes “truth” as more like “the color of the sky at noon” than “magnetism”.¹⁵ It functions like a disguised definite description, not of an object, but of a property. But unlike “the color of the sky at noon”, which picks out different properties in different environmental contexts, as it were, “truth”, on Wright’s view, refers to different properties in different *propositional domains*. That is, what determines which property “truth” refers to, and “is true” ascribes, is the domain the relevant proposition is a member of. Thus when saying it is true that acts of cruelty are wrong we ascribe one property; when saying that it is true that there is a book on the

¹³ “Minimalism, Deflationism, Pragmatism and Pluralism”, in Lynch (ed.), *The Nature of Truth* 752.

¹⁴ Here I am ignoring certain difficulties about the difference between rigid and non-rigid predicates (or singular terms standing for properties); Such difficulties are not a barrier to the contrast drawn in the text, which requires only that there is a distinction between descriptive terms or concepts like “is the color of the sky at night” and non-descriptive ones, like “is magnetic”.

¹⁵ Reasonable, but not the only interpretation; given the evolution of Wright’s position on the subject over the last decade, there naturally will be others. Compare, for example, Pedersen’s treatment of the position; *The Monist*, 89:1 (2006), 103–119.

table, we ascribe another. Nonetheless, in both cases—as with “the color of the sky at noon”—we employ a single concept, even though what property we pick out with that concept differs.

In sum, we can say that for Wright: in any particular domain, “truth” correctly applies to whatever property is described by the platitudes in that domain.

So revealed, the position bears a striking similarity to a form of reductionism in the philosophy of mind championed by David Lewis and more recently, by Jaegwon Kim. Kim, for example, holds that ascriptions of “pain” are non-rigid definite descriptions that denote different properties in organisms. On his view, there is no fact about whether, e.g. *x* is in pain over and above whether *x* has some physical property *P*, and so, as Kim puts it, “there is no need to think of [pain] itself as a property in its own right”.¹⁶ There is no one *property* that all the states we describe as “states of being in pain” share in common; pain-talk reduces to talk about brain states.

Wright’s view appears parallel. It seems as if there is no fact of the matter whether a proposition is true over and above whether it is superwarranted or correctly representing. And so we are tempted to conclude that for Wright, there is no property that all the propositions we describe as “true propositions” share in common; truth-talk reduces to talk about superwarrant, or correct representation. According to Wright, there may be something in common between all true moral propositions and all true propositions about middle-sized dry goods. But there is no property in common between all and only true propositions period in virtue of which they are true.

Looked at in this light, Wright’s official position doesn’t seem capable of solving the problems facing SAP after all. Consider first the problem of mixed compounds. The fact that we can apply a univocal concept of truth to the proposition that murder is wrong and grass is green does not tell us what property of that

¹⁶ J. Kim, *Mind in a Physical World* (Cambridge, MA.: MIT Press, 1998), 104.

conjunction itself the truth concept denotes. It seems to denote neither superwarrant nor correspondence. So what property does “realize” the concept (to use Wright’s terminology) in this case? Likewise in the case of generalizations: granted that “Everything Socrates said was true” applies a univocal concept of truth, what property does that concept express?

Mixed inferences cause another sort of problem for Wright. Typically we understand valid inferences as truth-preserving. This is meant to explain why an argument—any argument—is valid. The explanation tells us something about the premises and the conclusion: namely that if the premises have the property of truth, so must the conclusion. But if “true” denotes superwarrant in the case of moral propositions and correct representation in other cases, then in an example like the one we considered earlier, there is no single property preserved by the inference. Of course, the *concept* of truth is carried through on Wright’s view—that is, the premises and conclusion can both be *described* as true. But this fact does not explain why the description is apt; consequently it does not explain why the argument is valid.¹⁷ And it is hard to see how it could, since again, on Wright’s view, there is no property preserved by a valid mixed inference.

This last point suggests a tempting reply on Wright’s behalf. Perhaps Wright could reply that all true propositions and beliefs *do* share a common property. They all share the property of “falling

¹⁷ N. Pedersen, in “What Can the Problem of Mixed Inferences Teach us about Alethic Pluralism?”, 103–17, has suggested the Wrightian pluralist might avoid these problems by taking “true” to denote a single “plurality” of propositions rather than a property. (We accomplish this by employing Boolos’s plural interpretation of second-order quantification: roughly, “the proposition that *p* is true” when taken this way is equivalent to “it_i is one of them_j”). Ingenious, certainly—but does this help? It is rather like saying that being true is like being a member of a special club: the special propositions club. We can give the membership list of the club: *P* is a member of the special propositions club just when *P* is a proposition that is superwarranted, or corresponding with fact or.... But what makes all these propositions on the club’s membership list members of the club? If we’ve ruled out that “true” denotes a property, then we can’t say: because they share something—a property in common. So we must presumably say that what determines that *T*₁...*T*_{*n*} are all in the club is nothing more than their falling under the concept. Consequently, it is unclear whether ground is gained over the view discussed in the text.

under the concept of truth”. To fall under the concept of truth, on Wright’s account is to “satisfy the platitudes”. Thus we might reason as follows:

1. If a proposition is true, then it has a property that satisfies the platitudes. (premise)
2. Any property that satisfies the platitudes must be distinct from warrant, must be possessed by asserted propositions, must have a truth-apt negation and so on. (from Wright’s platitudes)
3. If a proposition is true, it has some property that is distinct from warrant, possessed by asserted propositions etc. (from 1, 2)
4. Therefore, all true propositions have *the property of having some property that is distinct from warrant, possessed by asserted propositions, etc.*

The argument’s last step relies on the following reasoning: if every true proposition has some property that has certain features, then there is something that every true proposition has in common. And if they all have something in common, then they share a property. What property? The property of having some property that has the relevant features.

Such a principle, of course, is controversial. But even if we grant the principle encoded in this reasoning, will this suggestion help? No. Wright is barred from identifying even this wafer-thin property—*the property of having a property that satisfies the platitudes or falls under the concept of truth*—with the, or even “a” property of truth. For the property of having a property that falls under the descriptive concept of truth, doesn’t itself fall under that description. Again, that description consists, essentially, in a list of the platitudes that a property must satisfy. But is the property of having a property that, e.g., is distinct from warrant, possessed by asserted propositions, is objective, and so on a property with all those features? No. It is the property of having a property with those features. Hence a view like Wright’s which identifies truth with whatever property satisfies the platitudes in a particular domain

must hold that the second-order property of having a property that plays that role is *distinct* from truth: call it truth*. And this in turn makes it hard to see how reductive pluralism solves the problem of mixed inferences and associated problems. For while she can say that there is a property preserved by valid mixed inferences, that property is truth*, not truth.

A common objection to Kim-style reductionism in the philosophy of mind is that it implies that pain is not a real psychological kind. There is nothing in common, in other words, between the states we describe as pain-states in dogs and the states we describe as pain-states in humans. Consequently the view gives up the ability to appeal to pain as such in general psychological explanation. And this is a loss. For we do find it useful and informative to talk about pain as such in order to explain other things of psychological interest, such as fear, or anger. Similarly with Wright's view. It implies that truth is not any more a real kind than pain is for Kim. The only property shared by all and only true propositions is one that is not, by the lights of the theory itself, ascribed by our use of "true" or denoted by "truth". Consequently, there is no property we ascribe by "true" that can be appealed to in order to explain certain general facts about belief, or validity, or (as we'll discuss later) content. And given that we often do take ourselves to appeal to just such a property in explaining such things, this too seems to be a loss.

In effect, Wright's view is far more deflationary in spirit than it appears. It gives up on truth as a general explanatory property, one which we can use in quite general explanations of other phenomena. To those who already believe that truth has no general explanatory role to play—who believe that it does not figure in explaining anything else of interest such as belief, or content or meaning—this will come of no surprise. But then they will not have needed Wright's view to reach the conclusion. On the other hand, to those like myself, who see truth as at least a potentially valuable explanatory resource, Wright's view remains dissatisfying.

Some of the arguments I just raised against Wright also weigh against a proposal I've made elsewhere—namely, that the pluralist

takes truth itself to be a 2nd order multiply realizable property.¹⁸ On this view, truth just is what we suggested the reductive pluralist would have to call truth* above; to be true is to have a property of having a property that satisfies the truisms or platitudes.¹⁹ Put differently, and drawing on an obvious analogy with psychological functionalism, we can say that truth is the property of having a property that *realizes* the truth-role.

The analogy with functionalism, as we'll see shortly, is apt. But truth can't be a second-order property, for the reason just noted. That would imply that truth is the property of having some property that has certain features. But does the second-order property itself have those features? That is, it seems that we want to say that *truth itself* is objective and a goal of inquiry. But now is my belief's having the property of having a property that is a goal of inquiry a goal of inquiry? Well compare: suppose the color red is a second-order property: being red is having the property of having a property with certain features, such as a reflective variance. Does the property of having a property with a given reflective variance itself have that reflective variance? Not obviously; indeed, obviously not.

Finally, it won't do to go disjunctive either. It might initially seem tempting to identify truth with the single disjunctive property, as in:

Truth = the property of being either superwarranted or corresponding to the facts.

There are some advantages to this approach for someone attracted to something like Wright's pluralism. A proposition will have

¹⁸ First made in Lynch, "A Functionalist Theory of Truth", in his *The Nature of Truth*, see also M. P. Lynch, "Truth and Multiple Realizability", *Australasian Journal of Philosophy*, 82 (2004), 384–408. A similar objection to my previous way of understanding functionalism can be found in Horgan and Potrč, *Austere Realism* (Cambridge, MA: MIT Press, 2008), 109–12. Thanks also to L. Shapins and D. Edwards (conversation).

¹⁹ This standard formulation of second-order properties is due to Kim: F is a second-order property over some set of base properties B if, and only if, F is the property of having some property P in B such that P satisfies some condition C on members of B. See his *Mind in a Physical World*.

the disjunctive property only if it has one of the properties that make up its disjuncts, so the pluralist intuition seems satisfied. And yet there is a single property preserved in valid arguments.²⁰ But whatever its other merits, the view founders on the problem of mixed compounds. Suppose, for sake of argument, that moral propositions are true by virtue of being superwarranted and propositions about physical objects and their properties are true by virtue of corresponding to the facts. Now consider the proposition that roses are red and murder is wrong. This proposition is true, but it can hardly be true by virtue of having the disjunctive property described above. As just noted, that would require that it have one of the “disjunct properties”. But it is not clear which property that would be. The proposition that roses are red might correspond but it hardly follows that the proposition that murder is wrong and roses are red corresponds to any fact.²¹

4. Two Demands

The upshot of this chapter is twofold. First, if we wish to avoid the problems faced by SAP, we need a univocal concept of truth. But second, if we are going to make sense of the idea that there is more than one way for beliefs to be true, we need more than that, more, in other words, than conceptual unity between all true propositions. Rather, we also need to make sense of the following metaphysical principle:

Truth is One: there is a single property named by “truth” that all and only true propositions share.

²⁰ Nikolaj Pedersen suggested this point to me in conversation.

²¹ One might think that there are different kinds of truth in the Aristotelian sense of genus/species. But the distinction fails to capture the pluralist’s thought. For x to be a species of y is for x to possess a property that together with the property of being a y , makes it an x . Thus the traditional example: man is a species of animal because to be a man is to be an animal and to be rational. But this isn’t what is going on in the case of truth: propositions which correspond to the facts don’t possess a further property—other than corresponding—that together with the property of being true makes them correspond to the facts. Corresponding to the facts can’t be defined by conjoining truth with some other property.

Yet, as the discussion in Chapter two illustrated, it would be good to be able to make sense of the intuition that drives pluralism about truth, namely

Truth is Many: there is more than one way to be true.

In the next chapter, I will begin outlining a theory that meets both demands.

4

Truth as One and Many

We say that the one and many become identified by thought, and that now, as in time past, they run about together, in and out of every word which is uttered, and that this union of them will never cease, and is not now beginning, but is, as I believe, an everlasting quality of thought itself, which never grows old.

Plato, *Philebus*

I. An Analogy

The idea that there can be one property that takes many forms is hardly unknown in philosophy. It is commonplace in the philosophy of mind in particular, where the thought that mental properties are multiply realizable is a mainstay. One mental property—pain, for example—might be realized by certain neural properties in humans, by very different neural properties in other animals, and by possibly still different properties in Martians. Whether an organism is in pain is dependent on it having one of these other physical properties: having one of those properties is what makes it have the relevant psychological property. Properties that supervene on a distinct kind of properties in this way are said to be multiply realizable.

Whether multiply realizability will ultimately serve to explain the nature of a property like pain, or the mind in general, something similar appears well suited for explaining the idea that motivates

pluralism—that truth is many—while still retaining the idea that truth is one. For it seems possible to treat truth, like pain, as a supervenient property: that whether a belief-content—a proposition—is true is determined, at least in most cases, by its having *some other property*. And this is compatible with saying that there is more than one property that could do this. That is, it is compatible with the thought—to put it roughly—that if a particular kind of proposition has the property M, it has the property of truth, but it is possible that there are propositions of another kind which are true but lack M. In other words, truth is a single higher-level property whose instantiations across kinds of propositions are determined by a class of other, numerically distinct properties. If so, then we have a straightforward way of meeting both of the demands canvassed above. What is many are the properties that intuitively make or determine that a proposition true; what is one is truth itself.

This chapter begins to develop this idea.

2. A Functional Theory of Truth

If there is more than one property that could determine that propositions are true, how would we identify which properties those are? This a natural question. And its answer is already implicit in our truisms. Recall that we identified the following three principles as historically prominent folk preconceptions we share about truth:

Objectivity: The belief that p is true if, and only if, with respect to the belief that p, things are as they are believed to be.

Norm of Belief: It is prima facie correct to believe that p if and only if the proposition that p is true.

End of Inquiry: Other things being equal, true beliefs are a worthy goal of inquiry.

Traditional theories of truth see these truisms—and, as we noted in Chapter 1, the truisms that are entailed by these three—as

giving us the nominal essence of truth. They believe that truth's real essence underlies the truisms and explains them. Accordingly, such theories try to reduce truth to some other property which they see as constituting truth's real essence.

But instead of seeing our truisms as fodder for reduction, why not see them as telling us what the property of truth does so to speak—about its function? Think of the Objectivity truism for instance, not as something that needs explaining by something more fundamental, but as revealing an aspect of truth's *functional role*. Roughly speaking, “portraying things as they are” is part of the “truth-role” in that it is something that propositions which have the property of truth do. Likewise with *Norm of Belief*: making a proposition correct to believe is just part of truth's job. In short, the thought is that the truisms tell us that true propositions are those that have a property that has a certain function in our cognitive economy. If we think of the truisms in this way then a natural answer to our question above is revealed: *the properties that can determine that propositions are true are those that play the truth-role*.

To play a functional role, in the sense intended here, is to satisfy a job description, one which picks out certain features possessed by anyone who has the job. Writing a job description involves listing the tasks any one who has that job must do, and specifying how that job relates to others in the immediate economic vicinity. We define the job in terms of its place in a larger network of jobs, all of which are understood in relation to each other, and by weighting some aspects of the job as more important or crucial than others. In the philosophy of mind, so-called “analytic” functionalists take job descriptions for mental properties to be given by our implicit folk beliefs about those properties. In the case of a property like pain, these include truisms like “the threat of pain causes fear” and “if you are in pain, you may say ‘ouch’” and “if you are hit in the head, you will probably be in pain” and so on. These platitudes tell us that a property plays the pain-role when it is related to certain other mental, behavioral, and experiential properties of an organism.

Likewise, we can take the truth-role to be carved out by our core folk truisms about truth. These truisms form a theoretical structure of sorts—one which illustrates the relationships between true propositions and propositions with various other properties such as warrant, belief, correctness, and so on. These features, as in the parallel case of functional properties in the philosophy of mind, will not be primarily causal in nature, but quasi-logical and explanatory. But the basic suggestion in both cases is the same: the unique relations that truth bears to other properties suffice to pin it down by jointly specifying the truth-role.

I suggested above that the idea that truth is many is best cashed out as the view that there is a plurality of properties that can make propositions or beliefs true. We've now seen that there is an easy way to identify which properties those are. There are some properties that play the truth-role and in virtue of that fact, make propositions that have them true. This allows us to give truth-conditions for the application of the truth concept itself as follows:

(F) $(\forall x)$ x is true if, and only if, x has a property that plays the truth-role.

Moreover, we've identified the truth-role itself in terms of the core truisms that make up our folk conception of truth. Earlier, we noted that it is plausible that a theory isn't about truth unless it incorporates the core truisms. The present suggestion is that we go one step further and say that our core folk truisms pick out certain relational features—call them the *truish features*—which a property must have in order to play the truth-role. This allows us to unpack the metaphor of “playing a role” Where P is a proposition, and T is a property,

T plays the truth-role if, and only if: P is T if, and only if, where P is believed, things are as they are believed to be; other things being equal, it is a worthy goal of inquiry to believe P if P is T ; it is correct to believe P if and only if P is T ...

Here the “...” reminds us that, in addition to our three historically prominent core truisms, there will be additional truish features specified by those truisms which—together with trivial non-truth theoretic facts—are epistemically underwritten by the above three. We discussed some obvious candidates for other core truisms of this sort in Chapter 1. Since we are here interested in laying out the structure of the theory, we can remain neutral on exactly what these may turn out to be.

Call this the *functionalist theory of truth*. So far, we’ve said two things about the theory. First, we’ve specified the conditions under which propositions are true. A proposition is true just when it has a property that plays the truth-role. Second, we’ve said what it is for a property to play the truth-role. A property plays the truth-role when it has the truish features specified by the truisms. What we haven’t said is what truth itself is.

Our discussion in the last chapter tells us what we *don’t* want to say. We don’t want to say, with Wright, that “true” functions like a definite description, serving to pick out different properties in different domains. Applied to truth, and conjoined with the supposition that there is more than one property that plays the truth-role, this suggestion results in the implausible view that there is more than one property denoted by “truth”. Neither, we’ve seen, should we say that truth is just the property of having a property that plays the truth-role. This secures for us a single property of truth, but at the expense of making it a property that is no longer described by our truisms. And given that our folk concept of truth is the concept of a property that satisfies those truisms, that would be a less than ideal result.¹

So neither of these options is what we want if we are to meet the two demands canvassed at the end of the last chapter—if we are to admit that truth is both one and many. Fortunately, a more natural option drops right out of our suggestion that we understand truth functionally. In Chapter 1, we noted that it is reasonable

¹ Contra, as noted before, my claims in Lynch 2001 and 2004.

to think that our truisms about truth give us the nominal essence of truth. More traditional theories then look elsewhere for truth's real essence. But there is another option: why not take the truish features as constituting, at least in part, truth's real essence as well as its nominal essence?

Here is a simple way of unpacking this thought that drops out of the above functionalist approach to understanding truth. Properties can have their features accidentally or essentially. *Being a color* is an essential feature of the property *being red*, but *being Tom's favorite color* is only an accidental feature of *being red*. Functional properties are defined by their functional role; that is, by the sum of their relational features. Those features can therefore be thought to be essential to it. Thus, the functionalist, like the monist, can claim that there is a single property and concept of truth. The property *being true* (or the property of truth) is the property that has the truish features essentially or which plays the truth-role *as such*. Constraining ourselves to our three favorite core truisms, we can say that it is the property that is, essentially, had by beliefs just when things are as they are believed to be; had by beliefs at the end of inquiry and which makes propositions correct to believe.² And the concept of truth is the concept of just that property.

Yet this approach also allows the functionalist to claim that truth is, as it were, *immanent* in ontologically distinct properties. Let us say that where property F is immanent in or *manifested by* property M, it is a priori that F's essential features are a subset of M's features.³ Since it is a priori that every property's essential features are a subset of its own features, every property manifests itself. So manifestation, like identity, is reflexive. But unlike identity, it is non-symmetric. Where M and F are ontologically

² Conversations with Douglas Edwards helped me get clearer on the point in the above two paragraphs.

³ That is, it is a priori knowable that F's essential features are a subset of M's features. The basis of the priori knowability is left open—it will depend on the particular case. In most cases relevant here, it will be a consequence of our concepts of F and M.

distinct properties—individuated by non-identical sets of essential features and relations—and M manifests F, F does not thereby manifest M.

The manifestation relation is similar to the determinable/determinate relation. It is a priori that the essential features of redness, whatever they are, are a subset of the features of being scarlet. Consequently, if one understands that something is scarlet one has all one needs to understand that it is red. But according to the traditional distinction, determinables cannot determine themselves, so the relations are distinct.⁴ We might say that determinable properties are one type of immanent property.

There is more to say about manifestation, but let's pause and see how it might apply to truth, using a toy example. Should a property of a proposition such as *its corresponding to reality* manifest truth, it will be a priori that the truth features are a subset of that property's features. This of course requires further elaboration, but the basic idea emerging here is this: properties distinct from truth may nonetheless manifest it by including truth's features among their own features.

Manifestation is a rational relation. When one property manifests another, it is a priori knowable that the one "includes" (so to speak) the other. So it is arguably not the relation that holds between psychological properties like pain and their neural correlates. Merely understanding that some organism is in some particular neural state is not sufficient for understanding it is in pain.⁵ But the cases are parallel in a significant further respect. Just

⁴ Determinable properties are actually distinct from the immanent functional properties discussed here in at least two other ways. First, determinants generally differ from one another along some linear ordering. Second, determinants of a determinable mutually detest one another, to paraphrase Armstrong. That is, nothing that is scarlet at some point and time can be crimson at that same point and time. Of course, there are different ways to understand what it is to be a determinable. And some uses of the term would group the two relations as the same. See for example, S. Shoemaker's use of the term in his "Realization and Mental Causation", in C. Gillett and B. Lower (eds.), *Physicalism and its Discontents* (Cambridge: Cambridge University Press, 2001), 78 and 80. For a recent interesting discussion of determinables, see E. Funkhouser's "The Determinable-Determinate Relation", *Notis*, 40: 3 (2006), 548–69.

⁵ The present account is influenced by Shoemaker, "Realization and Mental Causation". See also S. Yablo, "Mental Causation", *Philosophical Review*, 101 (1992), 245–80

as the psychological functionalist will claim that which physical property *realizes* pain in a given organism is determined by facts about the organism, the alethic functionalist will claim that where another property manifests truth for a particular proposition which property does so will depend on facts about that proposition. In particular, it will depend on (i) the proposition's logical structure; and (ii) the domain of inquiry to which the proposition belongs.

The first condition, as we'll see in more detail in the next chapter, is familiar from traditional correspondence views, according to which the only sort of propositions that correspond to facts are atomic propositions. The truth of compound propositions is understood to—minimally—supervene on the truth of their atomic components.

The second condition is parallel to the thought, familiar from the philosophy of mind, that whether a given neural property realizes pain depends on the kind of organism whose neural property it is. (Thus a neural property that realizes pain in one kind of organism may not realize it in another kind of organism even where members of the latter kind have that property). It is this second condition that captures the intuition—motivating pluralist accounts of truth—that which property manifests truth for a given proposition will depend partly on what that proposition is about.

Putting these two points together, the functionalist about truth can say that for atomic propositions, no further property (other than truth) manifests truth full stop. For atomic propositions, ontologically distinct manifestations of truth are manifestations *relative to a domain*. Moreover, the functionalist will take it that, at least for atomic propositions; a proposition is true because it has that further property that manifests truth for that domain. That is,

the functionalist will embrace, where P is an atomic proposition of some domain D,

Necessarily, P is true if, and only if, P has the property that manifests truth for propositions of D.

Just as a single neural property may play the pain-role for one kind of organism and not for another, a single property like superwarrant may play the truth-role for propositions of one domain but not for propositions of another. Manifestations of truth for atomic propositions are manifestations for a domain. Putting this thought together with our analysis of manifestation, we get (where P is an atomic proposition of domain D and has property M)

Necessarily, P has the property that manifests truth for propositions of D if, and only if, it is a priori that, when had by atomic propositions of domain D, the truth features are *a proper subset* of M's features.

I will say more about domains in a moment. For now, the important point is that an atomic proposition is true when it has the distinct *further* property that plays the truth-role—manifests truth—for *the domain of inquiry to which it belongs*. Not being true consists in lacking that property, either because there is no such property, in which case the content in question is neither true nor false, or because there is such a property, but the proposition in question fails to have it, in which case it is false.

So here, roughly and without some of the lingo, is the basic picture. Propositions about different subjects can be made true by distinct properties each of which plays the truth-role. Thus (atomic) propositions about the antics of the ordinary objects and properties of our daily life may be true because they represent those objects and properties in one of the senses canvassed in Chapter 2. For propositions of that kind, correct representation plays the truth-role and it is a priori that if a proposition correctly represents it will be true. For propositions of another sort, perhaps moral

propositions, superwarrant may be what plays the truth-role, or manifests truth. But such properties play the truth-role, if at all, only accidentally. That is, they may have the truish features, but only when they are possessed by atomic propositions of a certain domain. In contrast, truth itself has the truish features essentially, or plays the truth-role *as such*.⁶

So far, we've sketched only the most minimal outline of the functionalist view. But the basic picture already helps us understand how truth can be both many and one. Truth is many because different properties may manifest truth in distinct domains of inquiry. In those domains they have the truish features. Truth is one because there is a single property so manifested, and "truth" rigidly names that property.⁷ In all possible worlds and contexts where "truth" refers at all, it refers to the property that has the truish features essentially. It is the unique property that is, necessarily, objective, had by beliefs at the end of inquiry and which makes a proposition correct to believe. This allows us to say that "truth" will be a name or rigid designator of this property. In all possible worlds and contexts where "truth" refers at all, it refers to the property that plays the truth-role as such—the role picked out by the core truisms in the actual world.⁸

3. Manifestation and Domain Specificity

We've noted that which further property manifests truth for a given proposition depends (i) on that proposition's logical structure and

⁶ The functionalist theory identifies truth with a unique property with certain features. In other words, truth is that property that has certain properties, and only those properties, essentially. The fact that, like any other property, truth has properties does not make it a second-order property. Something has a second-order property just when *it* has a property of having some distinct property. This is distinct from saying that truth is a first-order property that (again, like any other first-order property) itself has certain other properties.

⁷ Likewise, "is true" will be a rigid predicator—a rigidified description ascribing a single property in all worlds.

⁸ Thus in worlds where the inhabitant's folk truisms describe features which are other than those described by the core truisms in the actual world, the property that necessarily has those features will not be truth.

(ii) the domain of inquiry to which it belongs. More needs to be said about both; in this section, I'll address the latter condition; the former will be the topic of the next chapter.

What makes a proposition a member of a particular domain? The obvious answer: the subject matter it is about.

The obvious answer may seem too obvious, and therefore unhelpful. So we can and should say something more. But before doing so, let me note that it isn't clear why we *shouldn't* be content with the obvious answer. After all, it should come of no surprise to anyone reading this book that any philosopher interested in the familiar puzzle described in the Introduction will take it for granted that we believe all sorts of different kinds of propositions: propositions about ethics, mathematics, about the sundries of everyday life. No one, presumably, will deny that these propositions concern not just different subjects, but *fundamentally* different subjects. And *any* philosopher who wishes to claim that we should treat propositions about these subject matters differently—for example, by saying that they aren't representational, or are all false—must have a way of distinguishing propositions of different types from each other. In short: it *is* obvious that the proposition that two and two are four is fundamentally different in kind than the proposition that torture is wrong. But making out what this difference consists in is hardly the sole job of the functionalist about truth.

Nor is the functionalist the only philosopher who will confront troublesome cases: cases where it will not be clear whether a given claim is about mathematics or aesthetics for example. (Consider: "The number seventeen is beautiful".) Suppose, like many philosophers, one is inclined to realism about mathematics, but expressivism or subjectivism about aesthetic judgments. Confronting these issues is not just a task for the functionalist, but for anyone interested in addressing the correspondence puzzle.

That said, it would be nice to say at least something more about what, at least in general, distinguishes propositions of one

domain from another. A natural answer is: the kind of concepts (moral, legal, mathematical) that compose the proposition in question.

This of course raises the question of what makes a kind of concept the kind of concept it is—a substantive question in the theory of concepts which I will not address here. But we can say this. One kind of concept differs from another by virtue of (a) its relation to, and (b) the character of, the properties that kind of concept is a concept of. This should be uncontroversial. Insofar as it makes sense to distinguish our thought about morality as different from our thought about physics, (and surely it does make sense) that distinction must ultimately derive from differences between the concepts that compose such thoughts, and therefore the differences between the properties, if any, those concepts are concepts of.

Importantly, this does not dictate in advance how one should understand the concepts of any particular propositional kind. Properly, the present theory remains neutral on substantive questions about how to understand the various concepts and propositions of philosophical interest. In particular, it remains neutral on which domains will, for example, satisfy Responsiveness and which will be epistemically constrained. Thus it remains neutral on which domains manifest truth by correspondence and which manifest it by superwarrant. This is fitting: a theory of truth does not solve all metaphysical problems. Rather, it is the other way around: *your overall metaphysical views* determine what manifests truth for any particular domain.

These considerations suggest that belonging to a particular domain is a feature an atomic proposition at least, has in virtue of being the sort of proposition it is. Propositions are the kind of propositions they are essentially; therefore, belonging to a particular domain is an essential fact about an atomic proposition. Moreover, our theory implies that there can be at best only one property that manifests truth for every domain. Two consequences follow from these points when we consider them together. First, the pluralist's

looked-for possibility of *variable manifestation*: propositions of different kinds have distinct properties that manifest the truth-role. Yet second, *domain specificity*: no atomic proposition is a member of more than one propositional domain.

One might think domain specificity too specific by far. Holding that propositions are not sentences can allay this worry. The same *sentence* can be uttered in different discourses or conversational contexts. Sentences are not joiners by nature. No sentence has an intrinsic relation to any subject matter, for any sentence can be used to express different propositions. So in uttering the sentence: ‘the healthcare system is broken’ I may in fact be expressing any number of propositions, including that medical care is unjustly distributed, that it is too expensive, that the quality of care has diminished and so on.

Uncontroversially, one often needs to determine exactly what has been said on some occasion before one can determine whether what was said was true or false. The functionalist merely adds that what manifests the truth (or falsity) of my utterance depends on what I said. Nonetheless, it may turn out, that due to the intrinsic vagueness and fluidity in the concepts I am employing, a claim, even suitably disambiguated, may express an atomic proposition that is neither determinately a member of one domain nor another. In such a case we shall say that its truth is similarly indeterminate.

Understanding the pluralist theory of truth in this way, as implying that simple propositions are domain-specific in virtue of the domain-specificity of the concepts of which they are composed, has three advantages. First, it explains why propositions can have distinct properties that manifest the truth-role—by being different kinds of propositions. Second, it is compatible with a natural way of understanding what makes a proposition the kind of proposition it is—or, alternatively, what explains the difference between propositional domains. If we take propositions instead to be domain-independent, then we can no longer say, as seems intuitive, that kinds of propositions are individuated by the distinct

types of propositions and concepts that compose them. We would have to provide a further theory of domains. Third, it undercuts the troublesome possibility that a single simple proposition (not just a sentence) might manifest truth in one domain but not in another, since no atomic proposition is a member of more than one domain.⁹

4. Is this a Theory of Truth?

The pluralist about truth begins with the intuition that there is more than one way for propositions to be true. In this chapter, I've suggested that this is best understood as the view that truth can be manifested in a plurality of properties, each of which can play the truth-role.

I've so far laid out the bare bones of the resulting functionalist theory of truth: our concept of truth is the concept of a functional property. This property is the property that plays the truth-role as such, or has the features definitive of that role essentially. This property is immanent in—or manifested by—other properties in virtue their playing the truth-role in particular domains. If we put pluralism and functionalism together the result is an over-all view of truth. Truth is an immanent functional property that is variably manifested.

In following chapters, I'll develop and defend this idea. But before we fill out the theory further, it is natural to pause and ask whether the theory counts as a theory of truth by the standards introduced in Chapter 1. There we said that a theory counts as a theory of truth only if it incorporates the truisms into the theory. The functionalist account just given defines truth in terms of those truisms. For a proposition to be true just is for it to have a property that makes it correct to believe, just is for it to be the sort of proposition that we aim at believing in inquiry, and is objective.

⁹ One could avoid this consequence by relativizing truth to domains, that is, by claiming that there is no single property of truth but only truth-in-D and truth-in-D₁ and so on. With domain specificity, however, we avoid the consequence *without* relativizing truth.

So given that the theory defines truth in terms of these truisms, it obviously incorporates them.¹⁰

So our account counts as a theory of *truth*. But we also said that to count as a *theory* of truth, as opposed to a mere chat about it, an account must explain the truisms or at least explain them away. But it might seem that the functional theory doesn't really count by this standard as a theory of truth. Traditional theories of truth, like the correspondence or superwarrant theories, explain our folk truisms about truth by appeal to features of true propositions that go beyond those picked out by the truisms themselves. But in contrast, the present theory, as (F) indicates, does not explain our folk truisms in this way. On the contrary, one might say, it apparently takes them at face value and uses them to explain what truth is. This seems to be a trivial explanation if any explanation at all.

In fact, the functional theory of truth does explain our truisms. But it does so in a two-step manner. The first step is indeed trivial. Take the truism that it is correct to believe what is true; according to functionalism this is so because true propositions have a particular immanent functional property, one which is defined in part as the property propositions have when they are correct to believe. The second step is not trivial. For functionalism also tells us that propositions are true only if they have a property that manifests truth. According to our definition of manifestation, a property manifests truth only if it has the "truish" features in some particular way. Consequently, depending on what property manifests truth for a particular proposition, we can say that what makes it correct to believe *that* proposition is that it has the property of superwarrant, or correspondence.

An important additional point here is this. Functionalism does not imply that the core truisms exhaust *all the features of truth*, or

¹⁰ Note that the theory is consistent with Warrant Independence. For that truism claims that some propositions can be warranted but not true and some true but not warranted. The functionalist theory tells us what truth is: the property that plays the truth-role as such. One can have that property but not be warranted, and warranted without having that property.

even all the essential features. For the fact, if it is a fact, that truth is a variably manifested property is surely a substantive, essential, fact about truth not found among our folk-theoretic beliefs about it. It would be silly to say that the grasp of the folk concept of truth requires a tacit understanding of a pluralist metaphysics of truth. If our analysis of the folk *concept* of truth is correct, then our ordinary concept of truth is the concept of a functional property. But if our pluralist *metaphysics* of truth is right, then this property is manifested in more than one way.¹¹ Since this is not a truism about truth, but a further fact about it, the theory that results from combining these ideas, the theory, in short, that truth is an immanent functional property that is variably manifested, is a substantive theory—one which tells us something new about truth. In short, we might say that while truth must have the features specified in its nominal essence, that nominal essence does not exhaust its real essence.

The point of this chapter has been to introduce the key elements of the functionalist theory of truth. The next two chapters further articulate and refine the view.

¹¹ It is a further, and to my ear, implausible claim that truth *must* be manifested by more than one distinct property. For whether truth is variably manifested depends on, among other things, the kinds of propositions we express. And nothing guarantees that we express the same kind of propositions in every possible world. Perhaps there are worlds where we do not express moral propositions.

5

Truth, Consequence, and the Universality of Reason

Logic takes care of itself; all we have to do is look to see how it does it.

Wittgenstein, *1914–1916 Notebooks*

Truth, I've claimed, can be profitably understood as an immanent functional property. Understanding truth in this way allows us to capture the core insight behind pluralism while retaining the thought that there is only one property of truth. What is one is truth; what is many are the ways in which that one property is manifested.

In the last chapter, I noted that the functionalist will take it that which property manifests truth for a given proposition depends not only on the subject matter the proposition is about, but on the logical structure of the proposition. But so far we've only defined truth for propositions with a simple logical structure. We've said that an atomic proposition of some domain is true if, and only if, it has the particular property that manifests true for propositions in that domain. But compound propositions may still seem to present a problem. In particular, mixed compounds: what property manifests truth for the proposition that *murder is wrong and electrons have negative charge*? Surely not some mixed "moral/physical" property?

The so-called "problem of mixed compounds" is not a mere technical issue. It is simply one face of a fundamental worry about

any theory of truth that tolerates pluralism—whether about truth itself or about the properties that manifest truth. That deeper worry concerns how we reason across domains of inquiry. Reason, by its nature, is universal in its scope—it allows us to combine propositions from different domains into more complex propositions, and to make inferences across different subjects—as when we draw moral conclusions from partly non-moral premises. But the very universality of reason raises an obvious question for the functionalist. If what it is to be true *varies* across domains, how can reason universally apply in every domain? If alethic functionalism is correct, how are we to understand the scope of reason?

1. Mixed Compounds and Plain Truth

The problem with the universality of reason has several faces. The two most striking concern what we earlier called *mixed inferences* and *mixed compounds*.¹ At first glance, functionalism would seem to provide answers to both. Strictly speaking, this is the case; but matters are apt to get a bit more complicated, as we'll see.

Consider first the problem of mixed inferences. In Chapter 3, we noted that this was a key problem for simple alethic pluralism. For if “true” was simply ambiguous, we were to understand inferences such as:

If hitting someone causes pain, then it is wrong
 Hitting someone does cause pain
 So it is wrong.

¹ For other discussions, see C. Tappolet, “Mixed Inferences: A Problem for Pluralism about Truth Predicates”, *Analysis*, 57 (1997), 209–10; M. Sainsbury, “Crispin Wright: Truth and Objectivity”, *Philosophy and Phenomenological Research*, 56 (1996), 899–904; M. P. Lynch, *Truth in Context* (Cambridge, MA: MIT Press, 1998), ch. 5. See also M. P. Lynch “A Functional Theory of Truth”, in his, *The Nature of Truth* (Cambridge, MA.: MIT Press, 2001), 723–50; and J. Dodd, “Recent Work on Truth”, *Philosophical Books*, 43 (2002), 279–91; N. Pedersen, “What Can the Problem of Mixed Inferences Teach us about Alethic Pluralism?”, *The Monist*, 89:1 (2006), 3–117; D. Edwards, “How to Solve the Problem of Mixed Conjunctions” *Analysis*, 68:2 (2008), 143–9; A. Cotnoir, “Generic Truth and Mixed Conjunctions: Some Alternatives”, *Analysis*, 69:2 (2009).

Typically we understand valid inferences as truth-preserving: if the premises are true, or have the property of truth, so must the conclusion. But if “true” means “superwarranted” in the case of moral propositions and “correct representation” in the case of propositions about the causal bases of our mental states, then there is no single property preserved in such inferences. But if we now understand truth as an immanent property, then a single truth property *is* preserved in such inferences, and moreover, we are able say why this is so by giving a general characterization of consequence. To wit: a valid inference is one where truth is preserved across its manifestations from the premises of the argument to its conclusion.²

The problem of mixed compounds, on the other hand, may seem to be more difficult. The functionalist theory maintains that every true proposition has a property that manifests truth or plays the truth-role. But then as noted above, how is the functionalist going to understand the truth of claims like *torture is wrong and the book is the table*?

Of course, even putting aside “mixed” examples, the truth of compound propositions is a general problem. Correspondence theories, for instance, have typically faced embarrassing questions about whether there are conjunctive or disjunctive facts to which conjunctive and disjunctive propositions correspond. Thus the functionalist, like the correspondence theorist, must have something to say about compound propositions.

An obvious tactic—and again one often adopted by correspondence accounts—is to appeal to a broadly recursive strategy. It is open to the functionalist, as it is to any theory of truth, to apply the theory in the first instance to atomic propositions, and then to understand the truth of a compound proposition in the standard

² In a related point: a blind generalization such as “Everything God believes is true” will mean just what it says: it is a generalization over propositions believed by God all of which are claimed to have the single property of truth. We can acknowledge that God, should he exist, believes all sorts of propositions. We can acknowledge that many of these propositions will have their truth manifested differently.

recursive way, namely as a truth-function of the atomic propositions of which it is composed. The truth of compound propositions is a logical consequence of the truth-values of their component parts, together with rules governing the use of the relevant connectives. Applied to our functionalist theory, we will say that the proposition that A & B is true because it is a truth-functional compound of conjuncts both of which do manifest truth; the proposition that A or B is true because it is a truth-functional compound of two disjuncts at least one of which manifests truth and so on.

At this point a natural question arises: do compound propositions have a property that manifests truth or not? If they do, what sort of property is it?

The functionalist has at least two general answers she might give to this question. The more traditional answer would be to claim that compound propositions are true without manifesting truth. The rationale behind this answer is the same as what drove logical atomism or most truth-making theories. Thus according to Russell there was no need to appeal to disjunctive or conjunctive facts; disjunctive or conjunctive propositions were true or false depending on the truth-value of their atomic parts.³ Or as the early Wittgenstein remarked,

My fundamental idea is that the “logical constants” do not represent. That the *logic* of facts cannot be represented.⁴

One way of reading this is that the only goal of a truth-theory is to tell us how the atomic propositions are true (or in Wittgenstein’s terms, represent). Recursion takes care of the rest.

Applied to the present theory, this implies that compound propositions are true or false only in a *derivative sense* by being truth-functional compounds of propositions which can manifest

³ The Philosophy of Logical Atomism, in his *Logic and Knowledge: Essays 1901–1950* (London: George Allen & Unwin, 1956).

⁴ L. Wittgenstein, *Tractatus Logico-Philosophicus*, trans. C. K. Ogden (London: Routledge and Kegan-Paul, 1922) sec. 4.6132.

truth. Thus we can say, for example, that a $\langle A \& B \rangle$ is derivatively true just when $\langle A \rangle$ manifests truth and $\langle B \rangle$ manifests truth.

This traditional sort of account, with its associated distinction between true and derivatively true propositions, has some merits. It accords with Wittgenstein's intuition that nothing out in the world makes a compound proposition true save the truth of its component parts. Nor is it unexpected. Just as some truth-maker theorists hold that, strictly speaking, there are no compound truth-makers, so the functionalist on this approach holds that there are no properties which manifest truth for compound propositions.

But the account also has some costs worth bearing in mind. One such cost, seldom noted in the parallel literature on correspondence, is that it makes our extensional definition of truth disjunctive. As such, it entails *revising* our original theory. Our original theory implies that

For all propositions P , P is true if and only if it has a property that manifests truth.

Rather than accepting this, however, we now replace it with the following characterization:

A proposition is true iff it manifests truth or is a derivatively true truth-functional compound proposition.

This is tantamount to saying that any proposition that is true is either an atomic proposition or a proposition logically derived from such propositions. Of course our understanding of the property and concept, remains the same. Truth is the property that necessarily has the features picked out by the core truisms about truth.

A second complication of the account is that it is wedded to a fully satisfactory completion of recursive analyses for all compounds. Some compounds, such as counterfactuals and subjunctives, are notoriously difficult to understand in this way. So while the use of

recursive analyses in truth-theory is familiar, the nature and extent of such analyses is a vexed issue.⁵

Fortunately, there is a second way to answer the problem of compound propositions which, while building on the intuitions that drive the above analyses, is both simpler and follows more directly from our earlier account. Let us begin by granting what seems obvious: that whether or not a recursive account of all compound propositions can be given, there is something right about the insight that guides such analyses. What's right about it could be captured by saying that all truth is grounded in a certain sense. There can be no change in the truth-value of a compound proposition without change in the truth-value of *some* atomic propositions. The truth-value of compounds supervenes on the truth-value of atomic propositions. Call this the weak grounding principle.

According to the theory that truth is an immanent functional property, a property M manifests truth just when it is a priori that the truth-features are a subset of the features and relations of M. Manifestation, so described, is a reflexive relation, since every set is a subset of itself. Thus all properties, including truth, self-manifest. When a proposition is true only by virtue of self-manifesting truth, we can say that the relevant proposition is *plainly true*. What makes a compound proposition plainly true? Given our weak grounding principle, compound propositions are plainly true if their truth-value is grounded. That is, if their truth-value supervenes on the truth-value of propositions which

⁵ The issue of their general applicability concerns, as noted, has to do with the question of how to handle expressions which aren't straightforwardly truth-functional; while the nature of the analyses will depend on, among other things, the type of quantification involved and how 'full-fledged' a recursive analyses is attempted—e.g. whether it appeals to the structural components of the relevant truth-bearers (Kirkham (1992), 139; David (1999)). As far as I can see, the alethic pluralist who uses such analyses is not committed to any particular answer to any of these questions, nor that these issues remain any thornier for her than for any of the other numerous theories that wish to appeal to some form of recursion. Nonetheless, it would be good to avoid them. For commentary on these issues see, M. David, *Correspondence and Disquotation* (Oxford: Oxford University Press, 1994), 117–23; S. Soames, *Understanding Truth* (Oxford: Oxford University Press, 1999), 86–92; R. Kirkham, *Theories of Truth* (Cambridge, MA: MIT Press, 1992), 139ff.

are either compound and grounded or atomic (and hence whose truth-value depends on having a property *other than truth* that manifests truth).⁶

So even if it turns out that a recursive analysis does not apply to every compound proposition (like subjunctive conditionals, for example) the functionalist can accept the weak grounding principle. Moreover, she will have independent motivation to do so. For she is already committed to (a) the thought that what's true depends on what is true in a particular way; and (b) via her account of propositional domains, to the idea that true atomic propositions have further properties like superwarrant that manifest truth. Consequently, it seems reasonable for her to hold that a compound proposition's truth is ultimately grounded on the truth-values of atomic propositions. So compound propositions, are true because they are plainly true.

Assuming that the weak grounding principle is accepted, the second approach seems preferable; it is simpler and entails no revision to the functionalist theory. Every proposition, even compound propositions are true because they manifest truth.

2. More than One Logic?

So on the face of it, the functionalist offers a way out to the alethic pluralist from the vexing problem of the universality of reason. But the solution, as so far developed, ignores an important—and pressing—possibility. That possibility is that even the sort of alethic pluralism allowed by functionalism may well bring logical pluralism in its wake. And as we shall see, this brings back the problem of the universality of reason with a vengeance.

Logical pluralism is the thesis that there is more than one logic governing our reasoning. Since logics can be individuated by their

⁶ Of course, any *atomic* proposition that is true will have two properties which manifest truth: truth, and whatever ontologically distinct property *manifests truth for that domain*. But according to our functionalist theory as laid out in the last chapter, an atomic proposition won't *be* true unless it has the property that manifests truth for that domain.

account of consequence, one can say that logical pluralism is the view that there is more than one relation of logical consequence, or validity. Intuitively, an argument is valid when its premises necessitate (in some sense) its conclusion. And as we've already noted, validity is usually defined in terms of truth: an argument is valid just when if the premises are true the conclusion is (or must be) true. So, we might ask, if there is more than one property that manifests truth, could this entail that there is more than one way for an argument to be valid?

Prior to grappling with this question, it will be helpful to say something about logical pluralism in its own right. Jc Beall and Greg Restall have recently argued for just such a position. They do so by defining validity by reference to what they call "cases", as so:

VALID: An argument is valid if and only if, in every case where the premises are true, so is the conclusion.⁷

They argue that this *minimal concept of validity*, however, is permissibly enrichable in more than one way, so long as the enrichment satisfies three platitudes about consequence: that it is a necessary relation, that it is a normative relation, and that it is a formal relation in at least some of the relevant senses of that term.⁸ Thus, for example, one might endorse:

CLASSICAL: An argument is valid if and only if in every possible world where the premises are true, so is the conclusion.

Here we take the "cases" referred to in VALID to be classically constrained possible worlds. An altogether different enrichment would be to take cases as *stages of inquiry*:

CONSTRUCTIVIST: An argument is valid if and only if at every possible stage of inquiry where the premises are true, so is the conclusion.

⁷ J. Beall and G. Restall, *Logical Pluralism* (Oxford: Oxford University Press, 2006), 27.

⁸ Beall and Restall, *Logical Pluralism*, 14–20.

Here stages of inquiry are understood as they are in our definition of superwarrant above. They are both extensible (additional information might always come in) and inclusive (the additional information is just that—additional; all successive stages of inquiry include the information warranted at prior stages). Again, as with our definition of superwarrant above, stages are potentially incomplete—a given stage of inquiry may neither warrant a claim nor its negation. Consequently, we lack warrant for thinking that the law of excluded middle holds in all stages, and likewise for double-negation elimination. Roughly speaking, we can say that CONSTRUCTIVIST is an intuitionistic logic.⁹

One does not need to adopt Beall and Restall's definition of validity in terms of cases, however, in order to understand logical pluralism. In particular, our definition of superwarrant allows us to state Beall and Restall's CONSTRUCTIVIST definition of consequence in terms of possible worlds:

CONSTRUCTIVIST*: An argument is valid if and only if at every possible world where the premises are superwarranted, so is the conclusion.¹⁰

Once again, the ensuing logic does not accept the law of excluded middle (LEM) since it is consistent with the notion of superwarrant that there is some proposition P such that we will have no reason to believe that either it or its negation are superwarranted, and therefore no reason to believe that LEM holds of all propositions. Likewise with the semantic principle of Bivalence.

Now suppose that CONSTRUCTIVIST* were to govern our reasoning in some domains, but CLASSICAL governs our reasoning in the rest. If so, then we would, presumably be

⁹ Obviously CLASSICAL and CONSTRUCTIVIST are only stand-ins for full logics; moreover there are other logics one might want to consider (and Beall and Restall do consider).

¹⁰ Here, we are not, as will become clear below, assuming that possible worlds are *complete* sets of propositions.

committed to what we might call *domain-specific logical pluralism* (or DLP): distinct domains of inquiry would be governed by different logics.

In any event, two relevant questions arise about DLP and its relation to alethic pluralism. First, does any type of alethic pluralism entail this sort of logical pluralism, and in particular, does the functionalist version sketched above entail it? Second, whether or not the functionalist is *committed* to DLP, how would her view be affected should she, for whatever reason, think it is true?

Let's take these questions in order. Strictly speaking, there appears to be no direct argument from any alethic pluralism to DLP. This is because if consequence is, roughly, a matter of truth-preservation, then just because there is more than one way for a proposition to be true does not mean that there must be more than one way for truth to be *preserved* from one proposition to another. Truth is one thing; truth-preservation is another. So it might be that there is only one way for truth to be preserved from premises to a conclusion even if there is more than one kind of truth to be preserved. Suppose, for example, that truth is plural but no kind of truth requires a revision of the classical laws of logic. If so, then CLASSICAL would be the one and only consequence relation.¹¹ The conclusion is the same when we restrict our attention to functionalism. Validity could be a functionally defined relation and still only be manifested in one way.

So there is no direct route from functionalism about truth to DLP. Nonetheless there is, it seems, an *indirect* argument. It is indirect because it involves several additional, if plausible assumptions: namely, first, that truth is variably manifested, and second, that one of the properties that can manifest truth is

¹¹ The case is perhaps not as clear when the question is whether logical pluralism entails alethic pluralism. Beall and Restall deny it (see (2006), 100); Stephen Read (in conversation) suggested otherwise. I remain neutral on the issue as of this writing.

superwarrant. Earlier we argued, with Wright, that superwarrant is a candidate for manifesting truth in any domain which meets (along with other constraints):

EC: If P, then it is feasible to have warrant for believing P.

Assume that there is some domain of which this principle is true. Grant that superwarrant or some similar property plays the truth-role for the propositions of that domain. If stages are defined as CONSTRUCTIVIST* defines them above, then, intuitively, the law of excluded middle cannot be known to hold for that domain. The intuitive case rests on—the admittedly plausible sounding assumption—that there is no guarantee that inquiry is complete at any stage. If so, then there may be propositions of the domain for which no warrant either for or against is ever available—even in principle. Consequently, there may be some proposition P in the relevant domain such that we are not warranted in holding: superwarrant P or superwarrant \sim P. And if we aren't warranted in accepting that, then neither, presumably, are we warranted in accepting that LEM holds for every proposition. Consequently we are not warranted in including LEM into our logic for that domain.¹² CONSTRUCTIVIST* is intended to allow for this. So if superwarrant manifests truth in a given evidentially constrained domain, then the consequence relation in that domain will be better construed as CONSTRUCTIVIST* and not CLASSICAL.

Thus the *indirect* route from alethic functionalism to DLP: if there is more than one way to manifest truth, and some of the manifesting properties are epistemically defined properties like superwarrant, and some not, then different domains will admit of different manifestations of the consequence-relation. And this means, among other things, that argument forms that are valid in some domains may not be so in others. All this of course, assumes

¹² This is not to say that LEM is false; see Crispin Wright, "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism", *Mind*, 110 (2001), 45–98.

that there is more than one way to play the truth-role. If there is not, then there may still be more than one consequence relation, but this will presumably be motivated by other things than a view about the nature of truth.¹³

The indirect argument for the conclusion that alethic functionalism involves a commitment to DLP can be resisted along a number of fronts. The two most obvious routes are first, to deny that truth is variably manifested; or second, to claim that our truisms about truth themselves constrain the logic that govern any domain. The former route is to deny pluralism. The latter is open to any one who wishes to claim, for example, that among the *core* truisms that demarcate the truth-role are foundational principles of classical logic, including paradigmatically, Bivalence. If this principle is a core truism, it picks out an essential feature of truth. It will therefore be a necessary truth that every proposition is either true or false.¹⁴

In any event, both of these routes are consistent with the functionalist theory of truth presented in this essay, and the latter is consistent with pluralism as well. Thus—to stress the point again—nothing forces the alethic pluralist to be a logical pluralist, domain-specific or other-wise.

But if we assume for the moment that none of these routes for resisting the indirect argument *is* taken, the question remains as to how an admission of domain-specific logical pluralism will affect the alethic pluralist. The *general* upshot, again, is this. In domains where propositions are made true by an epistemically constrained property like superwarrant, and where it is plausible that not every proposition will be either true or false, then the logic in that domain will be best modeled by CONSTRUCTIVIST*.

¹³ Compare Beall and Restall, *Logical Pluralism*, 100–1.

¹⁴ A third route to resisting the indirect argument would be to hold alethic pluralism would be to insist that CONSTRUCTIVIST* (or something like it) is the universal logic. We will return to this more complicated option below.

Two specific upshots are these. First, anyone who claims that distinct domains of discourse are governed by different logics must say something about mixed compounds, such as:

- (1) The cat is on the mat and torture is wrong.
- (2) Torture is wrong or grass is green.

Second, the functionalist/logical pluralist must also say something about mixed inferences, e.g.

MIX: Torture is wrong or grass is purple; grass is not purple, so torture is wrong.

Earlier we noted that for functionalists about truth, there is only one property of truth, even if that property can be variably manifested. Thus it is that single property that is preserved is a valid inference like MIX. Likewise, compounds like (1) or (2) are not true in some special “mixed” sense of “true”, nor are they true in virtue of some special mixed property of truth. They are true because (a) they self-manifest truth; and (b) their truth-value is grounded. So the functionalist *qua* functionalist, has no particular problem with either mixed inferences or mixed compounds. These issues arise again only when logical pluralism is on the table.

Propositions like (1) and (2) above are both “plainly” true, like any true compound, mixed or not. But as our reasoning above indicates, the fact that all compounds are true in the same sense does not solve the problems generated by DLP. And if, as we are assuming for the moment, the functionalist is likely to be a logical pluralist, then the problems of DLP are problems for the functionalist.

Let’s look at these problems in a bit more detail. Consider any conjunction or disjunction where one component is a proposition from a domain where truth is manifested by superwarrant, and the other component is from a domain where truth is manifested by some form of correspondence. Suppose, for sake

of argument, that the moral domain admits of a non-classical (CONSTRUCTIVIST*) consequence relation as defined above. The question will then be what to say about mixed compounds like the following:

- (3) Grass is green and Sophie's choice is morally right.
- (4) Grass is green or Sophie's choice is morally right.

Sophie, as in the book by William Styron, is forced by the Nazi's to choose which of her two young children will live and which will die. Suppose that our moral theory tells us that it is simply not decidable whether her resulting decision is the right one in the sense that we will never have warrant for or against it. If so, then we have no warrant for thinking either that it is superwarranted that her decision is right or superwarranted that her decision is not right. Thus we will not accept that Sophie's choice is morally right or it is not; nor, if truth is superwarrant, will we accept as true or false that

- (S) Sophie's choice is morally right.

What then do we say about the value of (3) and (4) (assuming that "grass is green" expresses a true proposition)?

Linked to this problem is the question of what to do about mixed inferences, such as MIX or

- NIX: If it is not the case that Sophie's choice is morally right, then grass is not green. But grass is green; so Sophie's choice is morally right.

The example is a toy one. But the possibility it raises is real. We often do infer across domains, and we often wish to infer moral conclusions from premises that include some non-moral considerations. Yet which logic should we use for the evaluating such inferences? One which, like CLASSICAL, counts NIX as valid? If so, then we are in the unfortunate situation of countenancing as valid an argument that fails to preserve truth—since the conclusion

(S), is (let us continue to imagine) counted as neither true nor false in the moral domain.¹⁵

It should be stressed that the above problem does not arise solely for those who've come to DLP via truth pluralism. The issue of how to deal with mixed inference and compounds is an issue for any logical pluralist who takes it that distinct logics operate in different domains of discourse. Fortunately for both sorts of pluralists, there appears to be a single solution available to logical pluralists for the problem of mixed inferences and the problem of mixed compounds.¹⁶

The solution has two parts. First, the advocate of DLP, being a pluralist after all, will take it that *within a domain*, what qualifies as the governing logic will be determined by what manifests truth in that domain. Thus, where the propositions that compose a compound are all from a single domain, and the premises and conclusion of a given inference are all from a single domain, the appropriate logic will be that which governs the atomic propositions from that domain. So if the domain of propositions about physical objects is governed by CLASSICAL, then all inferences and compounds within that domain will be governed by CLASSICAL as well. This is all just a natural consequence of being a logical pluralist, as we've understood the position.

The second, and central, aspect of the solution is to endorse a principle of *logical modesty*. Two recommend themselves. The first can be summarized as follows. Let's say that a compound's *weakest member* is the atomic proposition whose domain has the weakest logic relative to the logics in play so to speak. (That is, relative to the other logics governing the domains of the other atomic propositions composing the compound in question.)

¹⁵ Moreover, the argument will count as invalid by the lights of CONSTRUCTIVIST*, since it implicitly employs double-negation-elimination twice, the final use of which would be traditionally disallowed by such logics. So we might as well ask: is the argument valid or invalid?

¹⁶ I was aided here by discussion with Aaron Cotnoir and Nikolaj Pedersen.

Likewise, an argument's *weakest member* is that premise whose logic is the weakest in play. The weakest logic in play is that which has the fewest logical truths or which sanctions the fewest valid inferences.

MODEST: where a compound proposition or inference contains propositions from distinct domains, the default governing logic is that of the compound or inference's weakest member.

Thus in the case of NIX above, the first premise is the weakest member because it is a compound whose weakest member is a moral claim. And (we've been imagining) the moral domain is best modeled by a logic that is weaker than CLASSICAL, that is, CONSTRUCTIVIST*. Hence, according to MODEST, the inference itself is governed by CONSTRUCTIVIST*, and is therefore not valid.

Likewise in the case of (3) and (4). Here again (S) is the weakest member in the defined sense, and thus in each case the governing logic will be the logic of (S). How this pans out for the truth-value of either (3) or (4) depends, of course, on how one understands that logic. But for sake of illustration, suppose that we take the logic governing (S) to be a CONSTRUCTIVIST* logic according to which LEM and Bivalence are not guaranteed to hold of every proposition. If so, then (3) will not be true and (4) will be true. (3) will not be true because its second conjunct, (S) will not be true (or false); (4) will be true because its first disjunct is true.

The MODEST principle implies that, under this assumption, CONSTRUCTIVIST* is the default logic for mixed compounds and mixed inferences. In saying this, we imply that the logic appropriate for plain truth is CONSTRUCTIVIST* unless

- (a) The propositions that compose the compound or inference are all from a single domain.

And

- (b) The property that manifests truth in that domain is epistemically unconstrained.

This gives us a rule, in effect, for deciding which logic is appropriate for any given kind of proposition.

Of course the solution just offered depends on whether a motivation can be supplied by the advocate of DLP for MODEST. Interestingly, one drops right out of a previous commitment to pluralism about *truth*. The truth pluralist will take it that some kinds of propositions will have their truth manifested by correspondence, while others will have their truth manifested by superwarrant. Likewise, the domain-specific logical pluralist will take it that CLASSICAL may well govern our inferences in our thought about certain domains, perhaps those of the natural sciences and mathematics. But in other domains, such as the moral or aesthetic domains, it will not. In the latter sort of domains, we might well take it that claims like (S) are indeterminate in truth-value. Why? Our pluralist twice-over has an obvious answer: because such indeterminacy *is enforced by the property that manifests truth in the relevant domain*. And this is what we should expect, if as is natural, we take it that the underlying nature of truth for a domain dictates what logic holds for that domain. Consequently, given this direction of explanation, we will not want to endorse as valid any inference that would violate our truth pluralism. That is, we don't want to count as valid any inference that would require that propositions like (S) be either true or false. To put it still another way, when reasoning across domains, logical caution is in order: we want to limit the number of logical truths that we endorse, so as to respect those domains which, by virtue of the property that plays the truth-role within them, enforce less logical laws than others.

So alethic pluralism supplies a motivation for MODEST. And that is an implication worth flagging all by itself. For it is only by adopting MODEST, I've argued, that the domain-specific logical pluralist can deal with the problems for their position that arise from the universality of reason. And that means that, anyone attracted to DLP—anyone, for example, who thinks that different logics might govern our reasoning about morality and

physics—would do well to take alethic pluralism very seriously as well.

One reason some may be uncomfortable with MODEST is that it depends on the assumption that the logics in question can be ordered, in the sense that the stronger logics are extensions of the weaker logics, or that all the models of the former are models of the latter. And of course not all logics are so strictly ordered. But one who embraces DLP need not accept that they are—just as doing so doesn't mean that one must embrace any old logic. Rather, what the domain-specific logical pluralist must accept is that the *domain-specific* logics are ordered. She need not accept that all logics are domain-specific. And it does not seem an unreasonable constraint on those logics that apply only to specific domains of inquiry that—in virtue of the content that composes that domain—they be capable of being ordered along a continuum of weaker to stronger.

A more concessive response to this worry is available however. This brings us to our second way of approaching the modesty condition. This second approach might be summarized as:

MODEST*: where a compound proposition or inference contains propositions from distinct domains, the default governing logic is that comprised by the intersection of the domain-specific logics in play.

The thought here might be put by saying that any domain-specific logics in play are partially ordered in the sense that there will always be a further logic that is comprised of their intersection. Like its cousin principle, MODEST* cautions a type of logical conservatism. Thus it will supply the same results for (3) and (4) as MODEST. This is because CLASSICAL contains CONSTRUCTIVIST* as a subset. Thus the intersection of the weaker CONSTRUCTIVIST* and CLASSICAL is itself CONSTRUCTIVST*. Moreover, MODEST* does not require that all domain-specific logics be ordered along a continuum of weaker to stronger. It is consistent, for example, with the

thought that some domain-specific logics are of equal strength.¹⁷ Consequently, MODEST* seems an attractive option.¹⁸

Finally, some might wonder whether there is a significant difference between adopting the modesty approach—in either of its guises—and claiming that while truth is plural, there is a single, weak logic, along the lines of CONSTRUCTIVIST*. Therefore, why not say that CONSTRUCTIVIST* is not only the default logic, it is the only logic, and that domains whose logic appears classical only do so because we are employing additional principles such as Bivalence which aren't part of the one true logic.

Alethic pluralism, at least in the functionalist guise I've presented here, is certainly consistent with this suggestion. But the suggestion comes at a price. According to the suggestion, CONSTRUCTIVIST* holds in all domains. Bivalence is not recognized as a logical principle by CONSTRUCTIVIST*. Therefore in every domain, bivalence is not recognized as a logical principle. Therefore in domains which, according to this suggestion, nonetheless *appear* classical—and therefore abide by bivalence—bivalence must be true for some *non-logical* reason. And one might wonder what that reason might be.

A full assessment of this suggestion, therefore, requires drawing the boundaries of logic, an issue well beyond the scope of the current essay. But one small point is worth making: it seems natural that if one domain allows some inferences as valid and another does not, they have different logics. And domains where bivalence holds will allow some inferences as valid that other domains (which don't sanction bivalence) will not. So the natural thought is that they have different logics. Now according to the present suggestion, it might be that some inferences are counted

¹⁷ As one might think would be the case if the two logics were duals of each other, as in the case of supervaluationism and subvaluationism. See D. Hyde, "Pleading Classicism", *Mind*, 108 (1999), 733–5 for discussion of this possibility.

¹⁸ Of course, it requires the assumption that all domain-specific logics (although not all logics) be commensurable. But that seems reasonable in any event, given the fact that we do in fact reason across domains—that is, given the universality of reason.

as valid in a given domain not because the logic counts them as so, but because there is an additional metaphysical assumption that, together with the logic, allows them to count as so. But that just seems to mean that principles which function like logical principles are not logical principles, and again, one might wonder why that would be.

The alethic functionalist, even the functionalist who believes that truth is variably manifested, is not *required* to endorse domain-specific logical pluralism. But it is likely that she will. And if she does, she is wise to be modest—a virtue in philosophy as well as in life.¹⁹

¹⁹ Thanks to Marcus Rossbeig, Patrick Greenough and Aaron Cotnoir for helpful advice and discussion about the central issues of this chapter.

6

Deflationism and Explanation

Everything should be made as simple as possible, but not one bit simpler.

Albert Einstein

I. Deflationism

To see truth as an immanent functional property is to reject traditional theories, which reduce truth to a single nonfunctional property like correspondence or coherence. In part, this is because the functionalist theory, unlike its more traditional counterparts, allows for truth to be variably manifested. As such, it avoids what we earlier called the scope problem.

Today, however, many—perhaps most—philosophers who work on truth for a living are apt to reject traditional theories for a different reason, namely because they ascribe to one version or other of deflationism about truth. Hence a natural question facing the functionalist theory is how it compares to deflationism. The question is all the more pressing because, as we'll see, the views are alike in some respects. Nonetheless, they are more unlike than alike; and indeed, alethic functionalism retains many of the benefits of deflationism while lacking some of its problems.

In order to make this case, I first need to say what deflationism is. This is easier said than done, since the term “deflationism” is a lot like “naturalism”. It covers a host of separate views, many of which are motivated by distinct concerns and differ in the details.

Nonetheless, much like the various naturalisms in philosophy, the various deflationary theories of truth share a few core tenets.

First, deflationists hold that the concept of truth is merely a logical device. They generally base this claim on the fact that we are inclined to a priori infer the proposition that snow is white from the proposition that it is true that snow is white and vice versa. As a reflection of this fact, deflationists typically give pride of place in their account of truth to some form of the equivalence or T-schema:

TS: $\langle p \rangle$ is true if and only if p .

Thus Paul Horwich, for example, holds that our grasp of the concept of truth consists in our inclination to accept the instances of TS.¹ Yet the concept one so grasps, Horwich thinks, is merely a device for generalization; it allows us to generalize over potentially infinite strings of claims. Instead of saying that Tom says that grass is green and grass is green; and Tom says that roses are red and roses are red... and so on for everything Tom says, I can employ the concept of truth and simply point out that everything Tom says is true. If I had a mind big enough to encompass infinite conjunctions, I wouldn't need the concept of truth. But being human I do. The truth concept is an instrument that allows us to overcome our natural cognitive shortcomings. And that is all that it is.²

Second, deflationists share a metaphysical view: truth has no nature. Early deflationists sometimes put this by saying that the concept does not denote a property. But more contemporary deflationists wisely avoid this claim. Instead, they allow that the truth concept does express a property—in the same sense that the concepts of existence or identity express either a property or relation.³ Such properties, we might say, are *metaphysically transparent*

¹ P. Horwich, *Truth*, 2nd edn. (Oxford: Oxford University Press, 1998).

² C. Hill's theory offers a variation: truth "reduces to substitutional quantification", which he describes as "a logical device". See his *Thought and World* (Cambridge: Cambridge University Press, 2002), 23.

³ H. Field, "Critical Notice: Paul Horwich's *Truth*", *Philosophy of Science*, 59 (1992), 322; P. Horwich, *op cit*.

or pleonastic properties. Metaphysically transparent properties have no underlying nature that isn't revealed in our grasp of the concept; grasping the relevant concept tells us the whole essence of the property.⁴ Another example of such a property would be the property *being a conjunction*. Our concept of conjunction is the concept of a proposition of the form A & B, such that it is true just when both conjuncts—A, B—are true. That is all there is to a proposition's being a conjunction, or having the relevant property. One knows *all* the facts about *being a conjunction* simply by grasping the concept of conjunction. Contemporary deflationists can be understood as arguing that truth is similar in this respect. Like the property *being a conjunction*, the truth-property is metaphysically transparent, in that we know all the essential facts about it—its real essence, as Locke might have said—just by grasping the concept of truth.

So far then: our concept of truth is a logical device; our grasp of that concept is revealed in our grasp of the instances of the T-schema; and all there is to the property of truth is revealed in our grasp of that concept. From this, the deflationist concludes that TS and its instances are the only facts about truth one needs to know in order to understand what truth is. Any other fact about truth can be deduced from them together with relevant non truth-theoretical facts.⁵ No further metaphysical investigation is needed to tell us anything about the property.

Deflationists often put this point by saying that the role of the truth concept is expressive not explanatory. Truth does not play a significant explanatory role.⁶ By a "significant" explanatory role, I mean that truth doesn't figure in any explanations except in its role as a useful generalization device. This is not to deny that the

⁴ See S. Schiffer, *The Things We Mean* (Oxford: Oxford University Press, 2003), 61 ff; McGinn, *Logical Properties* (Oxford: Oxford University Press, 2001); and N. Damnjanovic, "Deflationism and the Success Argument", *The Philosophical Quarterly*, 58 (2005), 53–67.

⁵ M. David, "Minimalism and the Facts about Truth", in R. Schantz (ed.), *What is Truth?* (Berlin: DeGruyter, 2002), calls this the adequacy thesis: Horwich endorses the principle; see *Truth*, 136.

⁶ Horwich, *Truth*; M. Williams, "On Some Critics of Deflationism", and R. Brandom, "Explanatory vs. Expressive Deflationism about Truth" both in Schantz (ed.), *What is Truth?*

truth concept can't figure in explanations at all. It can, just so long as the concept in question is the deflationary one—that is, its role is limited to acting as a device for making generalizations over potentially infinite strings of propositions. But for deflationists, the fact that a proposition has the property of truth can't be an essential, primitive part of an explanation for some other phenomena. In part, this is because an appeal to truth in this way would be an excellent reason to think that truth does have a nature, or is a substantive property worthy of further investigation. For if truth does play a significant explanatory role in philosophical explanation, then it follows that there are facts about truth that go beyond what is revealed by our grasp of the instances of the T-schema.

These three points—that the concept of truth is a mere logical device, that the property of truth is a metaphysically transparent property, and that truth plays no significant explanatory role—jointly constitute the core of the deflationary position as I'll be understanding it here. Beyond this point, differences between various deflationary views emerge. Some views, for example, scorn propositions, preferring to talk about sentence truth, while others prefer utterances.⁷ Deflationists also differ over how to *justify* our commitment to the instances of TS: for some deflationary views, TS and its instances are simply the consequence of the fact that the concept of truth is *strongly semantically* transparent: $\langle p \rangle$ and $\langle\langle p \rangle$ is true \rangle are, in non-opaque contexts, synonymous, inter-substitutable or “cognitively equivalent”.⁸ While others deny that the instances of TS need justification.⁹ They instead claim that the non-paradoxical instances of TS are themselves epistemically and

⁷ W. V. Quine, *The Pursuit of Truth* (Cambridge, MA: Harvard University Press, 1990); H. Field, *Truth and the Absence of Fact* (Oxford: Oxford University Press, 2001).

⁸ Ramsey, probably erroneously, is sometimes thought to have held that $\langle p \rangle$ and \langle it is true that $p \rangle$ are synonymous; for Ramsey's view see “The Nature of Truth”, in M. P. Lynch (ed.) *The Nature of Truth* (Cambridge, MA: MIT Press, 2001). Jc Beall takes the two to be “intersubstitutable”, while Field holds them to be cognitively equivalent; see his “Transparent Disquotationalism”, in Jc Beall and B. Armour-Garb (eds.), *Deflationism and Paradox* (Oxford: Oxford University Press, 2005), and Field, *Truth and the Absence of Fact*.

⁹ Horwich, *Truth*.

explanatorily basic, and our grasp of the concept of truth is contained in our implicit acceptance of the non-paradoxical instances.

Whatever its specific form, deflationism is a very attractive view. Its most obvious benefit is its relative ontological simplicity. Since truth is metaphysically transparent, there is one less property we need to have a theory of. It captures the semantic appearances, as it were, without positing any mysterious relation like “correspondence” or “coherence”.¹⁰ But neither does it have the failings of pluralism. It is not committed to an ambiguity view of truth. Nonetheless, like pluralism, it completely avoids the scope problem. The scope problem only arises for views which specify some property P that all and only true propositions have which makes them true. Deflationists deny there is any such property; or if they do allow there is such a property, it will not be a property that will rule out any particular type of proposition from being true. It will be a metaphysically transparent property. For according to the deflationist, any proposition that is fit to figure in an instance of TS can be true.

Given these benefits, it is not surprising that deflationism has become as popular a view as it is. But deflationist theories also have considerable costs. Some of these costs concern problems associated with particular versions of deflationism.¹¹ Some of these concern how best to state some deflationists views, or even whether you can coherently state them all;¹² others are about whether some deflationists can account for generalizations about truth and involving truth.¹³

¹⁰ For further remarks of this sort, see the introduction to Jc Beall and B. Armour-Garb (eds.), *Deflationary Truth* (Chicago, IL: Open Court, 2004).

¹¹ Horwich’s minimalist theory faces the well-known generalization problem due to Gupta, “A Critique of Deflationism”, *Philosophical Topics*, 21 (1993), 57–81. Field’s pure disquotationalism faces problems with explaining our intuition that we can attribute truth to sentences we don’t understand. See S. Shapiro, “Truth and Proof: Through Thick and Thin”, *Journal of Philosophy*, 10 (1998), 493–521.

¹² Traditional disquotational theories face such problems, see, for example, M. David’s *Correspondence and Disquotation* (Oxford: Oxford University Press, 1994).

¹³ Horwich-style minimalism faces this problem, as first noted in A. Gupta’s “A Critique of Deflationism”.

Over and above specific worries about particular deflationary views, however, there are at least two objections to the entire approach—objections that cause me to want to pursue a different sort of view.

The first objection is that deflationary views may not even qualify as theories of truth at all. I've argued that a theory of truth is a theory of truth just in so far as it incorporates the core truisms. It is not clear that deflationary views can do so.

A typical worry in this neighborhood is that the deflationist can't pay sufficient homage to the so-called "correspondence intuition", or the thought that beliefs are true because of the way the world is.¹⁴ I am sympathetic; but it is not at all clear that this worry is justified. For what is truly platitudinous about the correspondence intuition is captured by Objectivity, or the thought that

The belief that *p* is true if and only if, with respect to the belief that *p*, things are as they are believed to be.

As we've seen, representationalists and antirepresentationalists alike can accommodate this truism. And not surprisingly, the deflationist can accommodate it too. Recall our earlier argument that we can warrant acceptance of instances of TS by way of Objectivity together with some conceptually linked principles about belief and truth. That is, together with:

With respect to the belief that *p*, things are as they are believed to be if and only if *p*.

And

It is propositions that are true or false.

But of course this argument might be flipped on its head. The deflationist might argue that Objectivity is really just a consequence of accepting instances of TS together with our associated facts about belief and truth.

¹⁴ See, for example, M. Dummett, *The Seas of Language* (Oxford: Clarendon Press, 1993), 52.

As it happens, I have serious doubts about whether this last claim is true. To my mind, Objectivity is clearly more explanatorily basic than TS, with its commitment to propositions. And TS, after all, is a schema—so it is difficult to see how it could be more explanatorily basic than an implicit generalization like Objectivity. But this point is hardly decisive, for either way accepting the instances of TS is perfectly consistent with accepting Objectivity.¹⁵

In my view, it is not Objectivity that spells trouble for deflationism, but Norm of Belief and End of Inquiry. For these truisms suggest that it is a distinct salient fact about truth that it has a normative dimension, not captured by the T-schema alone. Here is one way of making the point. Take the truism

Norm of Belief (NB): It is *prima facie* correct to believe $\langle p \rangle$ if and only if $\langle p \rangle$ is true.

Normative facts are those that are “fraught with ought” in Sellars’ phrase. And, arguably, any fact that implies ought-involving facts is a normative fact. It is trivial that, other things being equal, I ought to believe what is correct, and thus by (NB), other things being equal, I ought to believe what is true. Since any claim that implies oughts in this way is presumably normative, (NB) is a normative fact about truth.

Deflationists sometime respond that Norm of Belief is not really about truth at all.¹⁶ Their argument is that one can derive (NB) from instances of TS together with instances of the premise that

(B) It is *prima facie* correct to believe $\langle p \rangle$ if and only if p .

Thus, it is claimed, (NB) doesn’t really display a normative fact about truth. Rather, it simply illustrates how the concept of truth can be used as a logical device—in order to generalize over the instances of schemas like (B) for example.

¹⁵ See W. P. Alston, *A Realist Conception of Truth* (Ithaca, NY: Cornell University Press, 1996); Horwich, *Truth*, ch. 7, advances a slightly different view.

¹⁶ See, e.g. P. Horwich, “Norms of Truth and Meaning”, in Schantz (ed.) *What is Truth?*. See also M. McGrath, “Lynch on the Value of Truth”, *Philosophical Books*, 46 (2005), 302–10.

It should be noted that in order for this point to be persuasive, it must be convincing that instances of (B) are not about truth. But why think they aren't? It isn't sufficient to claim that (B) or its instances do not mention the word "true". Indeed, if (B) were said to be a good paraphrase of (NB), then we would have just as much reason to think (B) was about truth as we do for thinking (NB) is about truth. Good paraphrases carry their ontological commitments with them. (Otherwise they wouldn't be good paraphrases).¹⁷ Moreover, it isn't plausible to simply claim that (B) or (NB) is about belief and not about truth. Belief and truth are interrelated concepts—thus it is more plausible to say that (NB) or (B) tell us something about truth *and* belief. They tell us that belief's standard of correctness is truth and that truth is the standard of correctness of belief.

Even putting the above point aside, it is difficult to see why we would accept instances of (B) in the absence of already being committed to (NB). For the list of (B)'s instances is an infinite list of normative prescriptions: a list of little belief norms as it were: it is correct to believe snow is white iff snow is white, correct to believe roses are red iff roses are red and so on. Why should we accept each of these individual norms? Individual normative prescriptions are justified by general normative principles. Consider promising: it is correct to keep your promise to Tom for the same reason that it is correct to keep your promise to Bridget: because it is correct, other things being equal, to keep your promises. So too with truth: it is *prima facie* correct to believe that grass is green for the same reason it is correct to believe that snow is white: because it is *prima facie* correct to believe what is true. The *general normative principle*—(NB) in this case—is in the epistemic driver's seat. Consequently we are justified in accepting instances of (B) only in virtue of accepting instances of (NB). So even if—as I just argued was implausible—(B) isn't about truth, it can't

¹⁷ For an elaboration of this point, see W. P. Alston's classic, "Ontological Commitments", *Philosophical Studies*, 9 (1957), 8–17.

be used to derive (NB).¹⁸ And, given the close tie previously noted between (NB) and the truism that truth is a worthy goal of inquiry, it seems equally dubious that that normative principle could also be derived from TS alone. And this should hardly be surprising: the normative character of a complex concept like truth is usually thought to be irreducible to its purely descriptive character.

So one concern I have about deflationism is whether it can fully account for core truisms about truth. A less direct but, in my view, more fundamental worry, is this. Deflationists of all varieties must remove truth from the philosopher's toolbox. They must convince us to give up truth as an explanatory notion, by which to explain meaning, or belief or the success of our actions. They must do so because, as noted above, it is sufficient for P to be a real and distinct property if P is part of a significant and informative explanation for some phenomenon Q. When it is, we can say that P is an explanatory property. The thought is a familiar one: something is real if we need to postulate it in order to make sense of something else that is real. Thus if truth were part of an informative explanation of some other phenomenon of interest, we'd have good reason to think that there is more to say about it than the deflationist wishes to admit.

Naturally, deflationists are prone to see this as a good thing. Here's Paul Horwich, for example:

We have seen that many controversies—regarding, for example, scientific realism, meaning, vagueness, normative emotivism, and the foundations of logic—are standardly assumed to interact essentially with the nature of truth. To the extent that the notion of truth is clarified and its *independence* of these problems established they can be certain to receive clearer formulation and be more amenable to resolution.¹⁹

¹⁸ For further elaboration of this point, see my "Minimalism and the Value of Truth", *Philosophical Quarterly*, 54 (2004), 497–517, and *True to Life* (Cambridge, MA.: MIT Press, 2004). For responses to some criticisms, see my "Replies to Critics", *Philosophical Books*, 46 (2005), 331–42.

¹⁹ Horwich, *Truth*, 119.

I think that this is the wrong way around. Why should we think that construing truth—and these other problems—so as they are independent of each other should help us to get clearer on what those problems are about? They appear, at least partly, to be about truth. So wouldn't it be more straightforward, *more true to the problems as we find them*, to accept this and appeal to a theory of truth in order to help solve them? Moreover, why would we want to limit our resources? Explaining meaning, or content determination for example, is hard enough without requiring that we don't appeal to truth.

This point also bears on one of deflationism's benefits—ontological simplicity. A fact about theory building is that simplifying a theory in one way tends to complicate it in others. That is so here. To mention just one example which we will have reason to return to later: while deflationists (may) have a simpler ontology of truth, they do so at the cost of having a more complex theory of the nature of meaning and mental content, or no such theory at all. In particular they are barred from giving a truth-conditional account of the nature of meaning and content.

That deflationists must remove truth from our explanatory toolbox is my most fundamental misgiving about deflationism. Yet the view has significant benefits—chief among it avoids the problems of traditional theories. It would be good, therefore, if we could find a theory—a theory that had this benefit but allowed us to retain truth as an explanatory tool. Not surprisingly, I think alethic functionalism is that theory.

2. But What's the Difference?

One might wonder whether there is enough of a difference between functionalism and deflationism to warrant the above optimism. As it turns out, there is.

Without a doubt, there are some important similarities between the view defended in this book and deflationism. The views have some of the same benefits. Like deflationism, the present view

does not reduce truth to a single underlying structural property of propositions. As such, it also easily avoids the scope problem. Moreover, the deflationist and the functionalist share a particular theoretical methodology: they characterize truth by appeal to certain truisms or platitudes, including the T-schema. In particular, rather than trying to explain the truisms by appeal to another property, both views define the concept and property of truth in terms of the truisms. Hence both views are able to explain why it is that so many of our beliefs about truth seem platitudinous.

But the similarities between the views stop here. To begin with, deflationists typically take truth to be fully explicated by a single platitude, the T-schema (or its instances). This goes together with the core deflationary tenet that the concept of truth is merely a logical device. While functionalists admit the centrality of the T-schema in their theory of truth, they take truth—either the property or the concept—to be a more complicated affair. The simplest reason for this is that they see the T-schema as insufficient for a full characterization of the truth-role. According to the functionalist theory, truth is to be understood primarily in terms of its relations to other concepts, including *Objectivity*, *warrant*, *belief*, and *proposition*. Consequently, the functionalist can admit that the concept of truth is useful for generalizing over strings of propositions, but still maintain that this is not the only salient fact about the concept.

So functionalism about truth has a more complex account of the *concept* of truth than deflationism. Yet if alethic functionalism only differed from the typical deflationist view by virtue of its relative conceptual complexity, then one might well feel that the view avoids deflationism only in letter but not in spirit. Indeed, as I noted in Chapter 3, Crispin Wright's own pluralist position is for this reason closer to deflationism than it first might appear. For on Wright's view, there is nothing in common between all truths other than the honorific property of being correctly called true. Insofar as there are properties referred to by ascriptions of the truth-predicate, those properties differ from domain to domain,

because “is true” functions like a definite description. Such a view, I argued, has troubling consequences.

The theory of truth I’m defending in this essay is also *metaphysically* richer than deflationism. According to that theory, again, truth as such is understood as an immanent functional property, that is:

Truth = the property that plays the truth-role as such, or has the truish features essentially.

Early deflationists denied that truth is a property at all. Clearly this isn’t the position of the alethic functionalist. Other deflationists have argued that truth is not, in Horwich’s words, a ‘substantive property’. Substantive properties, according to Horwich, are properties that admit of a constitution theory of the form

X is true = x is F.²⁰

Alethic functionalism is precisely such a constitution theory. Thus by this standard, functionalism implies that truth *is* a substantive property. It is the property which plays the truth-role as such or has the truish features essentially—that is, having the various features picked out by our core folk truisms.

Above, I noted that the best way of characterizing deflationism’s metaphysical commitments is to say that according to deflationism, truth is a metaphysically transparent property. Metaphysically transparent properties have no underlying nature that isn’t revealed in our grasp of the concept; grasping the relevant concept tells us the whole, or real essence of the property. Therefore the fact that truth is not a metaphysically transparent property according to the present theory is a sharp distinction between it and deflationism. There are several reasons to believe this is so.

In general, a real distinction between a property and concept is merited whenever there are features essential to something’s being F which go beyond what can be known just by possessing the ordinary concept of an F. Whenever that is the case, something’s having the property *being F* cannot be a mere construction

²⁰ Horwich, *Truth*, 143.

out of the (ordinary) concept. Three reasons suggest that this is how matters stand for whether propositions have the immanent functional property of truth. Consequently, that property is not a mere conceptual projection, and therefore not a metaphysically transparent property.

The first and most important reason is also the simplest. Alethic functionalism, like deflationism, uses the core truisms to explain truth. It defines truth in terms of those core truisms: the property of truth is the property that has the truish features essentially. Note that this does not imply that the truish features exhaust *all the features or even all the essential features of truth*. For as I noted in the last chapter, the fact, if it is a fact, that truth is variably manifested is a feature of truth that is not revealed by our truisms. Moreover, if truth is an immanent property, then presumably it is essentially so. That is, it is essentially such that it could be variably manifested. But this too seems to be a substantive further fact about truth; it is not a core truism and hence while an essential feature of truth, not a feature that truth has in virtue of the concept of truth.

The over-all point here is of course familiar: there may be facts about a property—even essential facts about it—that go beyond our folk concept of that property. A common example is that there are facts about the property being water (namely that is necessarily identical to the property being H_2O) that go beyond our folk concept of water. Similarly, if our functionalist analysis of the folk concept of truth is correct, then our ordinary concept of truth is the concept of the property that plays the truth-role as such. If our pluralist *metaphysics* of truth is right, then this property is an immanent property that is variably manifested. So there are facts about the property that go beyond what is revealed in the concept. Hence truth is not a metaphysically transparent property if truth is an immanent functional property.

Second, there are facts about the extension of the concept of truth not included in the folk concept itself. According to alethic functionalism, an atomic proposition is true when it has the property that plays the truth-role as such. But *whether it has that property*

necessarily depends on whether it has another property that manifests truth. Should truth be variably manifested, the nature and character of the relevant manifesting property, and whether a given proposition has or lacks it, is determined in part by further facts about the propositional domain in question, just as what determines whether an organism is in pain is in part the sort of organism it is. These facts conjoin to necessarily determine whether something has the immanent property of truth. That is, they determine what propositions are and are not in the extension of the truth concept. But they go beyond the ordinary concept; mere possession of the concept of truth does not reveal them.

Third, in general, functionalists can allow that there are constraints on membership in a functional kind that go beyond what is revealed in the concept of that kind. For functional kinds like *heart*, there are broad and structural constraints on membership.²¹ The point of these constraints, so to speak, is to ensure that any realizer of the kind has certain properties that allow it to fulfill the function of pumping the blood. Hearts could not be made out of tissue paper. But not all such constraints will be known prior to investigation. Some may emerge only *after* we've compared some realizers of a supervenient kind with other quite different realizers. Doing so will often give us a better understanding of how they all differently realize the function in question, while at the same time help us to appreciate the constraints under which anything that does realize that function must operate. In such cases, that means there are facts that are relevant to whether something *can* have the relevant functional property that go beyond the features reflected in the concept.

In the present case, we can allow that there are constraints on what can have the property of truth by acknowledging constraints on what can count as a proposition which go beyond what is revealed in the concept of truth. The facts about such constraints

²¹ See L. Shapiro, "Multiple Realizations", *Journal of Philosophy*, 97 (2000), 635–54, for a discussion of this point with regard to multiple realization generally.

may remain unknown prior to investigation. Thus, consider the debate between those who believe there are spatially and temporally neutral propositions like:

It is raining

and those who think that there are no such propositions.²² According to the latter group, sentences which appear to express such propositions actually express more fully articulated propositions like:

It is raining in Storrs at 10 a.m. on September 18, 2006.

The functionalist can of course remain neutral in this and similar debates over the nature of propositions, while acknowledging that the outcome has impact on what can have or lack the property of truth. And since knowledge of such facts may exceed that which can be known about truth via the concept alone, they are not plausibly thought of as mere projections of the concept.

The functionalist theory of truth is clearly distinct from deflationism, even while retaining several of the latter view's benefits. Moreover, it has a further benefit that deflationism lacks. Deflationists remove truth from the explanatory toolkit. Their view implies that we cannot appeal to the nature of truth to help explain other items of philosophical importance—such as intentionality or mental content. If we are deflationists, we limit our options when addressing these other problems. The benefit of the functionalist theory, as I will argue more fully in the next chapter, is that it keeps the benefits of deflationism, but allows for the possibility that truth is an important part of explanations of other phenomena of interest.

²² Discussions include, D. Kaplan, "Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals", in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press, 1989); D. Lewis, "Index, Context and Content", reprinted in his *Papers in Philosophical Logic*, vol. 1. (Cambridge: Cambridge University Press, 1998).

3. Is the Immanent Property of Truth an Explanatory Property in its Own Right?

We first need to address a prior question. There is a well-known debate in the philosophy of mind over whether psychological properties, when understood as functional properties, could have any distinct explanatory relevance over and above their realizing properties. Those like Kim who believe that they could not suggest that this is sufficient to show that a functionalist theory of mind is incorrect (see Kim 1998). A parallel question obviously arises for the idea that truth is an immanent functional property. The parallel argument is that unless it can be shown that the immanent functional property of truth does distinct explanatory work over and above the properties that allegedly manifest it, then even for those sympathetic to pluralism, there is no reason to favor alethic functionalism over a more deflationary pluralism like Wright's. On that view, as I've understood in this book, "truth" functions like a description of whatever property plays the truth-role.

Presumably, it is sufficient for P to be distinct from Q if P, but not Q, is part of a significant and informative explanation for some phenomenon. When it is, we can say that P is an explanatory property. The other direction, which the above argument would seem to invoke, is less clearly true. That is, it is not clearly necessary for the reality of P that P be explanatory. There is nothing logically incoherent, at least, in a real but purely epiphenomenal property. Nonetheless, the following weaker principle does seem plausible: to the degree to which we come to believe that P does no explanatory work, to that degree we will find ourselves *doubting* its reality. Accordingly, if we think that P is real, we'd be well advised to silence doubters by illustrating how it does explanatory work.

In fact, there are two questions about explanatory work in the present context. First, there is the familiar question of whether "truth" *ever* names a property that does explanatory work—whether that property be one or many. It is this first question that

often divides deflationists from non-deflationists, as we saw above. Non-deflationists have cited a number of facts that truth helps to explain. Two common examples are:

- (i) The conditions under which some sentences are true explain their meaning.
- (ii) The conditions under which some beliefs are true explain their having the content they do.

The explanatory relevance of truth to meaning and content—should it exist—is quite intimate. In each case, the explanation in question is said to be constitutive: the meaning of a sentence, or the content of a belief, is constituted or necessarily determined by its truth-conditions.

In addition to such constitutive explanations, however, truth is sometimes said to be an explanatory property in a different sense—namely in that there is a counterfactual dependence between a belief's being true and some phenomenon. This thought is at the heart of what is sometimes called the "success argument".²³ According to one version of this argument, I would be less likely to succeed in fulfilling my desires if I lacked true beliefs about how to get what I want, and more likely to succeed if I possess true beliefs about how to get what I want. Beliefs which are true are those that have the property of truth. Hence having beliefs with that property helps to explain why I get what I want when I do.

The explanation here is not constitutive: no claim is being made to the effect that a belief's being true is an essential part of practical success. Rather the point is like that made in typical causal

²³ The argument was originally developed in Putnam, *Meaning and the Moral Sciences*, (London: Routledge & Kegan Paul, 1978) and by Field, "The Deflationary Conception of Truth", in G. MacDonald and C. Wright (eds.), *Fact Science and Morality* (Oxford: Basil Blackwell, 1986). N. Damnjanovic, "Deflationism and the Success Argument", *Philosophical Quarterly*, 55 (2005), 53–67 offers an excellent discussion of the typical deflationist response to the argument. His own diagnosis of why the typical response fails appeals to some of the same principles appealed to below in my response to the present objection about the explanatory power of the role property of truth.

explanations: x wouldn't be the case if y lacked the property of truth.

Deflationists typically respond to the success argument by claiming that it is just one more example of the fact that we use the concept of truth as a logical device for generalization. Their point is that "having true beliefs means that Bridget is more likely to get what Bridget wants" is equivalent, through the T-schema, to a long conjunction:

(W) If Bridget believes that using the water-fountain will get her what she wants and using the water fountain will get her what she wants, then she will be more likely to get what she wants, and if she believes that talking to Patty will get her what she wants and talking to Patty will get her what she wants, then she will be more likely to get what she wants, and so on.

The only role the concept of truth plays is in allowing us to express the long conjunction. But since I can explain the success of any particular action without appealing to my beliefs having the property of truth, we needn't conclude that having beliefs with that property plays any explanatory role.

As I will indicate shortly, I find this response unpersuasive. But for present purposes I will put to one side the question of *whether* truth is an explanatory property and assume that it is. For the more immediate question here is simply whether truth, *considered as an immanent functional property*, has a distinct explanatory role over and above the various lower-level properties that could be said to manifest it.

The first thing to notice is the difference between the present problem and the related problem in the philosophy of mind. Functionalists in the philosophy of mind have to explain, presumably, how the role property of e.g. pain can be *causally efficacious* over and above its realizing properties. The type of explanation we think matters for mental properties is causal explanation in this strict sense. We want to know how mental properties have any powers, over and above those enjoyed by their physical realizing properties, to "push and pull" other physical properties. Showing

this is thought to be difficult in the face of the plausible hypothesis of the causal closure of the physical world.²⁴

It is not clear that this is the issue here at all. When the deflationist denies that truth figures in explanations of e.g. linguistic meaning or mental representation, or in explanations of the norms of assertion and belief, she normally isn't denying that truth has causal efficacy. Nor need her opponent take himself to be making that claim. The type of explanations at issue here are not obviously causal. Even in the case of the success argument, the counterfactual dependence of successful action on true belief is not easily seen as itself depending on a brute causal interaction between a belief's possession of an abstract property of truth and a physical action (cf. Damjanovic, 2005). All causal relations support counterfactuals, but not all counterfactual dependencies need be supported by causal relations.

Thus the question before us is better put in terms of whether the functional property of truth does *explanatory* work. One way for a property to do explanatory work is for instances of the property to be causally efficacious. But there are other ways for a property to figure in a significant, informative explanation for some phenomena—appeals to the property may serve as “best explanations” for why the phenomena occurs, or serve as essential parts of a “unifying” explanation for various types of phenomena. Indeed, these latter forms of explanation are typical of the explanations that philosophers deal in. In philosophy, like in science, we are inclined to think that some abstract feature of the world exists (e.g. numbers, possible worlds, epistemic or moral properties) when we find that we cannot make sense of other, better understood phenomena without it.

One reason for thinking that truth might *not* do explanatory work over and above its manifestations would be an exclusion argument.²⁵ Here's how the argument might go. Suppose we want

²⁴ See e.g. J. Kim, *Mind in a Physical World* (Cambridge, MA.: MIT Press, 1998).

²⁵ Kim, *Mind in a Physical World*, and his *Supervenience and Mind: Selected Philosophical Essays* (Cambridge: Cambridge University Press, 1993).

to explain why I succeeded in getting my drink of water by (in part) an appeal to some property of my belief B about how to get that drink. Say that B causally maps (in the sense of principle (CC) from Chapter 2) some relevant objects, events, and properties. Wouldn't an appeal to that property be sufficient to explain why I managed to slake my thirst? And if it is, wouldn't that exclude B's *being true* from being an explanation of the same event? B's truth seems superfluous.

This sort of argument is not very persuasive when one of the properties in question is immanent in the other.²⁶ The explanatory work done by a property is presumably a function of its features. Yet if property X is immanent in property Y then the essential features of X are, in effect, *part of* the features of Y. So, if truth is immanent in correspondence in the sense outlined by a theory like (CC) then truist features are a proper subset of that property's features. But parts don't compete with their wholes for explanatory work. Consider a frequently cited example of Yablo's: suppose an industrial scale is set to ring a buzzer whenever an object on the scale weighs more than 10,000 pounds.²⁷ What explains the ringing of the buzzer, we could say, is the heaviness of the object on the scale. You drive your car, which weighs 10,001 pounds, onto the scale and the buzzer sounds. Your car is *barely heavy*—but does its bare heaviness exclude its being heavy from being explanatorily relevant? It is hard to see how it could, since being heavy is just part of what it is to be barely heavy. Similarly, if truth is immanent in a property like causal representation, then being true is just part of what it is to be causally representing. Likewise, if truth is immanent in superwarrant in some domains. Consequently, there is no question of explanatory exclusion of truth by its manifestations.

Indeed, in many contexts, an appeal to the property of truth *does more explanatory work* than appeal to one of its manifestations.

²⁶ The point was first made by S. Yablo, "Mental Causation", *The Philosophical Review*, 101:1 (1992), 245–80; and "Singling out Properties", *Philosophical Perspectives*, 9 (1995), 477–502.

²⁷ Yablo, "Mental Causation", 487.

One reason for thinking so is that an explanation for why such and such is the case is a good explanation when it is commensurate with such and such being the case. That is, when it explains such and such with a minimum of irrelevant information. And surely in many cases, the fact that my belief about how to get what I wanted was *true* will be the commensurate explanation. That is, appealing to that fact for why I got what I want will cut to the point with a minimum of fuss far more than an explanation than appeals to a particular metaphysical theory of my belief's structural properties.

Moreover, while carrying less *irrelevant* detail, an appeal to the immanent property may well carry *more* relevant information. Consider again the case where Bridget gets what she wants. Based on the *general principle* that one is more likely to do so when one has true beliefs about how to get what one wants, we explain that Bridget got what she wanted because she had true beliefs about how to get what she wanted. What information does this general explanation give us that a particular explanation that appeals to her belief's causal mapping would not? The explanation in terms of her belief's causal mapping tells us why in the actual world she gets what she wants. But the explanation in terms of the truth of her belief tells us that *in any relevantly similar situation, as in the actual situation*, Bridget's having a true belief about how to get what she wants will be a good explanation for why she gets what she wants. So an explanation of Bridget's beliefs in terms of their truth gives us *modal information that an explanation just in terms of the underlying manifesting property does not*. Since any explanation which supplies relevant information that some other explanation does not is *prima facie* a better explanation, it follows that explanations of practical success that appeal to truth are better explanations than explanations which appeal only to truth's manifestations.²⁸

²⁸ J. Jackson and P. Pettit would call these "program explanations". See their "Causation in the Philosophy of Mind", *Philosophy and Phenomenological Research*, 50 (1990), 195–214.

This latter point also demonstrates why the typical deflationary response to the so-called success argument fails. The deflationist, recall, will explain my practical success in getting what I want by appealing to reasoning like this: If I believe that talking to Bridget will get me what I want and talking to Bridget will get me what I want, then I will be more likely to get what I want. No need, they say, to appeal to the truth of my beliefs—I need only appeal to the ways in which the world is. And insofar as the truth predicate is used at all, it is used as a generalization device that allows us to avoid having to state long strings of conditionals like the above. But again, by explaining the success of my actions in terms of the *truth* of my belief, I implicitly convey new modal information: I convey the information that *other* true beliefs would have also brought about success had the world been different than it is—even if, in fact, it had been very different, different in ways I *cannot even imagine*. In making claims about truth's link to success, I am ranging over worlds and situations that I might not even understand, and hence would not recognize were one to include them in some long list of conditionals like (W).

A final point here is this. In general, an appeal to supervenient properties is justified whenever they are needed to explain commonalities. And we obviously do employ truth to explain such commonalities. We want to say, for example, that what is in common between all valid inferences is that they preserve truth, or that for any proposition, it is correct to believe it only if it is true. But such explanations do not come for free. In order to make them, we must reify over the supervenient, immanent property of truth. In response to this sort of point, some might think that I could sufficiently explain both the particular facts and commonalities mentioned above by appealing to the disjunctive property of being either causally isomorphic or superwarranted or.... M1 ... or.... M2.... etc. But any appeal to the entire disjunction of actual and possible truth manifestations will by its very nature bring in a host of information irrelevant to our particular

interests at hand. Such an explanation would not be commensurate; it is therefore inferior.

In this chapter, I have tried to accomplish two main tasks. First, I've argued that the functionalist theory of truth is not a deflationist theory. It offers a richer account of the concept of truth; more importantly, it implies that truth is not a metaphysically transparent property. This in turn implies that truth is no mere honorific property; it is a property that can do significant explanatory work in our philosophical theories of other phenomena of interest. Second, I've argued that the explanatory work that the immanent functional property can do is over and above the work that might be done by its various manifestations. In the next chapter, I expand this thought by examining some of the things that a functionalist theory of truth can be used to explain.

This page intentionally left blank

7

Expanding the View: Semantic Functionalism

You need to call to mind the differences between language games.

Wittgenstein, *Philosophical Investigations*

How one understands truth can affect how one understands a variety of other philosophically interesting concepts. In this chapter, I discuss some of the connections between the functionalist theory of truth and other areas of philosophical interest. In so doing, I hope to illustrate how the theory can, among other things, explain how the various domains of our thought can be both cognitively unified and yet semantically diverse.

I. Propositions

As I have presented alethic functionalism, it is propositions that are the primary bearers of truth and falsity. This naturally raises three questions. First, is the theory committed to propositions? Second, must it take propositions to be the primary bearers of truth and falsity? And, third, what account of propositions is being assumed? I'll address these questions in reverse order.

Propositions, as I use the term, are what stand to believing, disbelieving, asserting and denying as numbers stand to adding

and subtracting.¹ Propositions are what we believe, or the content of our beliefs. Propositions are the objects of our propositional attitudes.

The *prima facie* case for the existence of propositions is that they allow us to readily make sense of the fact that from

Terry believes that cats are cute

we can infer that

There is something that Terry believes.

Moreover, we can then add that what Terry believes is also what Michael believes and Tom doubts. In short, propositions allow us to treat “A believes that S” as being true when A stands in a believing relation to that to which “that S” refers. This in turn allows us to make sense of our common assumption that when you believe something is the case, there is something you believe.

Understood in this way, it is natural, from an ordinary language point of view, to take propositions as the primary bearers of truth and falsity. For propositions are what are believed, and (insofar as what I assert is what I believe) what is asserted. And it is what I assert or believe that we typically take to be true or false, not the act of believing or asserting. Likewise, it is not the sentence-token “grass is green” itself—the individual marks on the page—that we ordinarily believe is true.² Rather, it is what is said with that string of symbols. Likewise yet again, the noises I make when I speak are not what are true or false—it is what I say with those noises. So since what I say and believe is a proposition, it is plausible to take propositions to be the primary bearers of truth and falsity.

¹ I owe this particular turn of phrase to Peter van Inwagen.

² Neither is it the sentence-type either from the ordinary language point of view—for one sentence-type, e.g. “Joe put his loot in the bank” can be used to say quite different things; and that fact encourages us to see the bearer of truth or falsity to be what we use sentences to express—propositions.

I say “primary bearers” because I don’t mean to deny that we can ascribe truth to beliefs and utterances in a secondary way. That is, we certainly can say that

Terry’s belief is true.

But in saying this, we take ourselves to mean that *what* Terry believes is true. When we use such locutions, we are indirectly referring to the content of her belief—what I am calling a proposition—and ascribing truth to that. Consequently, it is still the proposition that is the primary bearer of truth in this case; the belief-state is true but only indirectly, via the truth of what is believed.

One might think that this commitment spells trouble for the pluralism component of alethic functionalism. For both possible manifestations of truth we’ve discussed—correspondence and superwarrant—traditionally take *beliefs* as the primary bearers of truth and falsity. In Chapter 2, however, we noted that both theories can accept propositions as the bearers of truth and falsity on the assumption (plausible on both views) that propositions—while they might exist mind-independently—can only be true or false only insofar as they are the possible objects of belief. In general, if a given property that manifests truth in some domain can only be born by believable propositions, then truth in that domain will only be had by believable propositions. But that doesn’t rule out that truth can be manifested by some other property that can be born by propositions that could not be believed.

Not surprisingly, the account of propositions I’ve just suggested is functionalist. The property *being a proposition* is a property something has in virtue of its satisfying a particular functional role, or doing a certain job. The relevant job description is marked out by a set of truisms, some of which I’ve just singled out as being at the core of our understanding of what a proposition is. Namely

Contents: Propositions are the objects of our propositional attitudes; they are *what* is believed and disbelieved.

Truth-bearers: Propositions are the (primary) bearers of truth and falsity.

Language-transcendent: The same proposition can be expressed by sentences in different languages and different sentences in the same language.

These “truisms” about propositions mark out what it is to play the proposition-role. There may be other core aspects of that role; some, for example, have argued that the following are also definitive of propositions:

Abstract entities: Propositions are independent of space and time.

Mind-independent: Propositions can exist independently of minds.³

While these latter principles are plausible, I prefer a thinner account of the proposition-role which is independent of either. On my view, the abstractness and mind-independence of propositions are metaphysical matters best left to debates over what fulfills the proposition-role, and are not essential to that role.

However thick, a functionalist theory of propositions is obviously consistent with a variety of proposals about what entities are marked out by these principles, or manifest the property of being a proposition. It is consistent, for example, with the idea that what fulfills that role are abstract structured entities whose component parts are the actual objects and properties that

³ S. Schiffer, for example, in *The Things We Mean* (Oxford: Oxford University Press, 2003) includes both of these principles in his list of core features of propositions. A familiar argument for why propositions must be mind-independent is that <roses are red> could be true even if there were no minds; hence that proposition could exist if there were no minds. It is unclear, however, whether such an argument is persuasive, since the premise is questionable—indeed, it is implicitly denied by traditional theories of the sort canvassed in Chapter 2—both of which deny that there can be unbelievable true propositions. Note moreover that our first three principles are consistent with an Aristotelean theory of propositions, according to which they are immanent in belief-states. Finally, Dummett points out that some abstract entities, like Chess, might nonetheless be dependent on minds in at least this sense: had there been no minds, Chess would not have existed. See Dummett, *Frege: Philosophy of Language* (London: Duckworth, 1973).

our beliefs are about (the Russellian account of propositions);⁴ or abstract structured entities composed of the concepts of the objects and properties beliefs are about (The Fregean account);⁵ or unstructured entities individuated by their possible world truth-conditions.⁶ It is even consistent with the idea that propositions are in Schiffer's phrase, "pleonastic", metaphysically transparent entities, about which there is nothing more to say than is already contained in the concept.⁷ Moreover, the theory is consistent with the possibility that the functional property of being a proposition is manifested in more of these ways, depending on the domain of inquiry in question.⁸ Of course, it is also consistent with its being the case that there is only one way that propositions are manifested.

In claiming that the conception of propositions discussed above is independent of these and similar questions about the nature of propositions, I am not suggesting that these latter debates are unimportant. Rather, I am suggesting that the immanent theorist is only committed to a broadly functionalist conception of propositions. A philosopher who embraces a pluralist account of truth like the one defended here is free, as is anyone, to adopt some particular theory of what manifests propositions; but her theory of truth does not itself dictate what theory of the nature of propositions she must adopt.⁹

⁴ See e.g. B. Russell, *The Principles of Mathematics* (New York: Norton, 1903).

⁵ Frege, "On Sense and Reference", *Frege Translations from the Philosophical Writings of Gottlob* in P. Geach and M. Black (eds.), (Oxford: Blackwell, 1952).

⁶ Stalnaker, *Inquiry* (Cambridge, MA.: MIT Press, 1984).

⁷ Schiffer, *The Things We Mean*.

⁸ Thus, a Russellian view might be thought to be more apt for propositions about middle-sized dry goods, for example, although I have no interest in defending that suggestion here.

⁹ Indeed, alethic functionalism is even more ecumenical than I have so far made it seem. As far as I can see, one could adopt its core tenets (that truth is plurally manifested) and yet hold that beliefs (in the act, not object sense) are the primary bearers of truth. The resulting view would not be the type of pluralism defended in this essay, but it would be a pluralist theory of truth just the same. Likewise with utterances or sentences. If, *pace* the considerations raised above, these items can be made out to be the primary bearers of truth and falsity, this would be no barrier to understanding truth in a pluralist way.

2. Semantic Functionalism

According to the theory of truth I've defended in this book, part of what it is to play the truth-role is to be the sort of property born by propositions. According to the functionalist theory of propositions just given, part of what it is to be a proposition is to be the sort of entity that bears truth. Some philosophers have taken this tight interconnection between truth and proposition to be an objection to the existence of propositions—and to more substantive theories of truth as well. As Wittgenstein notes,

We have

$P = \text{it is true that } p$

It is true that $p = p$

...it looks as if the definition—a proposition is whatever can be true or false—determined what a proposition was, by saying: what fits the concept or what the concept “true” fits, is a proposition. So it is as if we had a concept of truth, which we could use to determine what is and what is not a proposition ... but this is a bad picture. It is as if one were to say “the king in chess is the piece that one can check”. But this can mean only that in our game of chess we only check the king.¹⁰

Wittgenstein's point is that our concepts of truth and proposition are so intimately connected that it is not informative to say that propositions are what are true and false. The definition is trivial. Does the pluralist account of a proposition, coupled with the pluralist account of truth fall victim to this triviality worry?

In reply, note that the account of propositions offered here isn't simply that they are whatever bears truth. They are also language transcendent entities that are the object of our beliefs and assertions.

More importantly, the tight interdependence between truth and proposition is not surprising given our functionalist approach to

¹⁰ L. Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1953), sect. 136.

both properties. A commonly noted feature of functional properties is that they come as a package-deal. According to psychological functionalism, the property of being a belief has the functional role it does in part because it is related to other psychological properties like being a desire, and in turn, desires are intimately related to beliefs. That psychological properties are related in this way is simply a fall-out of psychological functionalism, not an objection to it. Moreover, no one takes this consequence to mitigate against the fact that functionalist explanations of mental phenomena fail to be informative. They are very informative—a functional account of a mental state like belief explains what it is by telling us about the role it plays in our mental economy—a role that relates it not just to the concept of a desire but also to action, experiential input, and truth.

The idea that *psychological properties* like belief and desire form a functionalist network of properties is familiar. So a natural conclusion to draw from the view that *truth* is also a functional property is that the various other semantic properties form a similar functionalist network. To draw this conclusion is to accept *semantic functionalism*, or the view that:

- Truth, meaning, denotation, content, proposition and other semantic concepts form an interrelated network of concepts/properties, all of which must be at least partly defined in terms of each other.
- The properties these concepts denote are subject to plural manifestation.

If semantic functionalism is true, then it would not be at all surprising to find that we must appeal to propositions to explain truth and truth to explain propositions. That is just what the theory would predict. And given their obviously interconnected nature, it should be expected that we learn about these properties as alethic functionalism would suggest, not one by one, but more or less all at once. Like learning to ride a bicycle, acquiring this skill requires the coordination of a whole set of tinier skills—balancing,

pedaling, steering, etc.—none of which are learned in isolation. You learn to master one as you come to master another, and vice versa. Understanding the properties of truth and proposition is similar in this respect.

To this we must add two caveats. First, the level at which we appeal to other properties to explain the nature of truth is at the level of understanding *truth as such*—not at the level of understanding what manifests truth in particular domains. Even if we accept semantic functionalism, therefore, we are not committed to saying that what we might call *deep* explanations of truth or content must exhibit the kind of circular interrelatedness exhibited by functional properties.

Second, nothing in alethic functionalism forces one to hold semantic functionalism. One can always hold that a property is functional in character without holding that the properties that are used to mark out its functional role are also functional. Nonetheless, semantic functionalism is a natural pairing with a theory that allows for pluralism like the functionalist theory of truth.

3. What Determines Content?

Truth is often thought to be connected to content, and in particular to what *determines* the contents of our beliefs. What implications does alethic functionalism have for our understanding of what determines content?

Strictly speaking: none. The theory is consistent with most of the dominant theories of content around today. But alethic pluralism does open the door to a new possibility, or perhaps better put—an old possibility that has recently been overlooked.

If, following the last section, we hold that the content of a belief is a proposition, the relevant question is not so much what content is as what determines content. Roughly speaking, there are three families of answers to this question.

Conceptual role theories hold that contents of mental states are completely determined by the role those states have in our thinking.

According to one influential theory of this sort, beliefs have the content they do in virtue of their component concepts, and those concepts have their respective content in virtue of their uses in inference, guiding behavior, and reactions to outside stimuli. In their pure form, such theories make content determination an *internal* matter—internal to the system of thought, or conceptual scheme, of the thinker.¹¹ Moreover, pure conceptual role theories are committed to some form of *holism* in that they take it that the content of a mental state is determined by its relation to all other mental states or at least some subset thereof.

A persistent worry for conceptual role theories, at least in their pure form, is that they don't account for the fact that *truth* seems to play a significant part in content determination. If I say XYZ is water when it isn't, I'm wrong. And the reason I'm wrong, one wants to say, is that our concept of water has the content, or *denotes*, H₂O not XYZ, and so my belief that some XYZ is water isn't *true*. Otherwise put, it seems that the belief that the glass contains water has the content <the glass contains water> at least in part because that belief is true if and only if the glass does contain water—which is to say that it contains H₂O. Hence the familiar externalist thought that the content of a belief is at least partly a matter of its truth-conditions, where these conditions are understood as external to our mental states and intentions. This is the key thought behind *truth-conditional theories* of content.

The exact structure of, and differences between, truth-conditional theories of content determination depends on all sorts of things of course, many of which I will not even begin to discuss

¹¹ The origins of what I'm calling pure conceptual role theories can be found in W. Sellars, "Some Reflections on Language Games", *Science Perception and Reality* (London: Routledge & Kegan Paul, 1963); and G. Harman, "Meaning and Semantics", in M. K. Munitz and P. K. Unger (eds.), *Semantics and Philosophy* (New York: NYU Press, 1974). Theories become less pure the less internalist they become—that is, the more they admit that roles are partly determined by relations to outside elements; see G. Harman, "(Nonsolipstic) Conceptual Role Semantics", in E. Lepore (ed.), *New Directions in Semantics* (London: Academic Press, 1987).

here. I'll focus instead on the simple approach that I earlier called representationalism.

Representationalism is often taken as both a theory of truth and as a theory of content. This is for the obvious reason that the theory of content (and meaning) is truth-conditional. It says, roughly, that a belief has the content it does in virtue of its having a particular truth-condition. And it explains that truth-condition in terms of (a) the belief's being structured by certain component concepts, and (b) the denotation of those concepts. Thus, we might say on this account that my belief B has the content that x is F in virtue of the fact that

B is true if and only if the object denoted by the concept x has the property denoted by the concept F.

This is, essentially, what I called in Chapter 2 REPRESENT. As I noted there, REPRESENT itself can be seen as something of a truism or platitude. Taken alone, it is consistent with deflationary theories of truth and reference.¹² More reductive accounts will go on to give a theory of what it is for concepts to denote. They might say, for example, that,

CAUSAL: <cat> denotes cats = cats cause, under appropriate conditions, mental tokenings of <cat>.¹³

Or,

TELEOLOGICAL: <cat> denotes cats = the biological function of <cat> is to be mentally tokened in presence of cats.¹⁴

¹² And of course with Davidsonian theories that, while not as deflationary as some, takes the concept of truth to be basic. See D. Davidson, *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984). Davidson himself of course has no truck with proposition and property talk.

¹³ For theories of this (broad) sort, see F. Dretske, *Knowledge and the Flow of Information* (Cambridge, MA.: MIT Press, 1981); D. W. Stampe, "Toward a Causal Theory of Linguistic Representation", in P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in Philosophy*, 2 (Minneapolis, MN: University of Minnesota Press, 1977); J. Fodor, *Psychosemantics* (Cambridge, MA.: MIT Press, 1987).

¹⁴ See R. Millikan, *Language, Thought and Other Biological Categories* (Cambridge, MA.: MIT Press, 1984); D. Papineau, *Reality and Representation* (Oxford: Blackwell, 1987).

Details differ, as do the problems faced by these sorts of theories. But as we noted in Chapter 2, a global problem for either CAUSAL or TELEOLOGICAL views is that the theories of truth they imply face a problem of scope. It is difficult to see how such theories can explain the truth of all of the beliefs we intuitively take to be true, including beliefs about norms and numbers. When used to develop a theory of content determination, this problem is magnified. The problem becomes not just whether the troublesome beliefs can be true, but how they can even have content—or even be beliefs at all.

Conceptual role and truth-conditional theories are two of the main kinds of theories of content. The third kind of theory is *verificationism*. On verificationist views, a belief has the content it does in virtue of the conditions under which it would be verified. Thus a belief's verification-conditions, rather than its truth-conditions, determine its content. Such theories are often motivated by the thought that an account of content should be able to explain how we can *grasp* such content and how we can *manifest* that grasp in our behavior.¹⁵ On a truth-conditional account, to grasp a proposition is to understand the conditions under which that proposition is true; but these conditions might transcend our ability to recognize whether they obtain. In such cases, our grasp of the relevant conditions may seem tenuous. Of course verificationist theories face this problem's mirror-image. For there seem to be some beliefs, such as the belief that there are an odd number of stars in the universe right now, which have content but which lack verification-conditions.

Nothing in the functionalist theory of truth requires its advocate to adopt any of these theories of content determination. But as noted above the theory does open a door closed to other approaches. That is, semantic functionalism, according to which core semantic concepts are (a) understood as comprising

¹⁵ M. Dummett, *The Seas of Language* (Oxford: Oxford University Press, 1996); C. Wright, *Realism, Meaning and Truth*, 2nd edn. (Oxford: Blackwell, 2001).

a functional network of concepts; and (b) the properties picked out by such concepts are understood as open to plural manifestation.

What this amounts to in the present context is this. Most philosophers find it intuitive that truth has at least some role in determining content. Yet if we are monists about truth, this implies that belief content, like truth, comes in only one kind. And this in turn, as we noted in Chapter 2, threatens to ignore the diversity of our thought. Semantic functionalism raises the following hope: perhaps we *could retain the key component of the truth-conditional account of content and yet maintain that there are different ways in which content is determined*. If so, then we might be able to retain some of the insights of two other views of content as well.

By the “the key component” here, I mean the modest claim that

TCD: A belief’s content is at least partly determined by its truth-conditions.

TCD invokes truth to help explain content determination. If the arguments of the last chapter are correct, the functionalist, unlike the deflationist, can maintain that truth does explanatory work. Consequently, she can adopt TCD. But given her theory of truth, she can also maintain that it is at least possible that the content determination relation is *itself open to multiple manifestation*.

Alethic functionalism, recall, holds that a (atomic) proposition is true only if it has the property M that manifests truth for its domain. And a property M manifests truth if, and only if, it is a priori that the truish features are a subset M’s features. Suppose, as I’ve argued is plausible, that superwarrant and correspondence can each be seen to manifest truth in domains that satisfy certain constraints. Accordingly, the conditions under which a belief in that domain is true will themselves depend on how truth is manifested in that domain. Where truth is manifested by superwarrant, the conditions under which a belief is true will be determined by the conditions under which it is superwarranted. Hence if TCD is

right, then the content of that belief will be partly determined by the conditions under which it is superwarranted. Likewise for domains where truth is manifested by correspondence. Another way to put this result is as follows: *Content is always at least partly determined by truth-conditions; but what manifests those conditions varies.*

Let's see how this might work in a particular case. Recall that in Chapter 2, I suggested that in some domains, it will be difficult to deny that our thoughts about Gs are responsive to the antics of the Gs themselves. In such domains, it will seem extremely plausible that

Responsiveness: Mental states with G-ish content are causally responsive to an external environment that contains Gs.

Consider the case of our thought about trees. It is difficult to deny, I would think, that Responsiveness holds for such thoughts. Human beings' beliefs about trees are affected by the fact that, among other things, we live in an environment where there are trees which we can perceive. Where Responsiveness holds for some kind of mental state, I earlier argued, then the relevant mental states should be understood as representational. And where states are representational, some causal correspondence theory of truth is at the very least a *contender* for the right account of what manifests truth. Thus, to take a simple version of the view (for illustration):

CC (Causal-correspondence): The belief that a is F is true if and only if the object causally mapped by <a> has the property causally mapped by <F>.

Here, to causally map an object is to stand in the appropriate causal relation that the advocate of CAUSAL believes is constitutive of denotation. My present point is simply this. Where CC is plausible, it is so partly *because* either CAUSAL or TELEOLOGICAL (or some other naturalistic, reductive view of denotation) is plausible. The views come together as a package. So, if TCD is true,

then where CC and e.g. CAUSAL is plausible, so too will be a particular theory of what manifests the relation of content determination.

Earlier I also maintained that if superwarrant is going to be a contender for being the property that manifests truth in some domain, at least three conditions would need to be satisfied:

Epistemic constraint. Where superwarrant is taken for truth in some domain, it must be in principle possible for someone at some time to have warrant for believing any given proposition of that domain.

Nonrepresentational content. The mental states we have toward propositions of the relevant domain are not subject to Responsiveness; hence they are not plausibly representational.

Norms of correctness. Given the truism that truth is the norm of judgment, if those beliefs are going to be truth-conditional, one must be able to make the case that said beliefs are nonetheless subject to normative constraint.

I have yet to argue that there are any domains that satisfy these constraints. But suppose there were. If so, then truth might be manifested by superwarrant, that is again:

Superwarrant: The belief that *p* is superwarranted if and only if the belief that *p* is warranted without defeat at some stage of inquiry and would remain so at every successive stage of inquiry.

Now suppose TCD is globally true. If it is, then in domains where truth is manifested by superwarrant, a given belief's content will be determined by its superwarrant-conditions. To put the point in deliberately Dummettian terms: content determination will be an antirealist matter in such domains.

4. Meaning and Truth-Conditions

One might think that a more radical type of functionalism about content is motivated by alethic functionalism. Why restrict, in

other words, what can manifest warrant to relations that satisfy TCD? That is, why should truth-conditions always be at least part of what determines content?

A reason for making the restriction can be seen when we turn our attention to the closely linked notion of meaning. The meanings of sentences are determined by the content of the concepts and beliefs such sentences can be used to express. Thus adopting the above view of mental content suggests a similar account of meaning. The most dominant view of meaning in empirically orientated philosophy of language is a broadly truth-conditional account. A benefit of semantic functionalism as I've been sketching it is that it is consistent with such an approach. Indeed, as I'll explain, it is consistent with it not just as an approach in formal semantics, but as a philosophical account of the nature of meaning.

There are two projects that can go under the heading of "truth-conditional semantics" or theories of meaning.¹⁶ One project involves giving what we can call a *semantic theory*. A semantic theory, among other things, distinguishes various semantic forms of linguistic items and explains how the meanings of complex linguistic expressions are built up out of the meanings of their component parts. Truth-conditional semantic theory has been reasonably successful at this task. It proceeds by assigning referents to words and then using principles about connective words to explain how to derive the referents of the more complex expressions from the referents of simpler expressions. Moreover, this method in turn allows us to offer explanations of other phenomena: it explains, for example, how my tacit knowledge of the referents of a finite number of words, along with tacit knowledge of how connective words work, allows me to understand the potentially infinite combinations of complex sounds that you could make when talking to me. And in turn, that helps to explain features of my psychology, such as the beliefs I form as a result of that understanding

¹⁶ N. Block, calls this "linguistic semantics" see for example, his "Conceptual Role Semantics"; Franklin Scott (ibid.) calls it "formal semantics".

and the behavior I engage in as a result of those beliefs. Formal truth-conditional semantics, in short, is a highly useful theory with important and testable empirical consequences.

Most theories of truth are compatible with the formal *machinery* of such a truth-conditional semantic theory.¹⁷ Thus even the deflationist is not barred from accepting Frege's thought that the truth of a sentence is determined by the reference of its parts, nor from accepting certain base clauses of a semantic theory, e.g. claims of the form

“cat” refers to cats

“mat” refers to mats

And ensuing statements of truth-conditions as

“The cat is on the mat” is true if and only if the cat is on the mat.

Taken as purely formal apparatus, a semantic theory allows us to use “true” and “refers” to map certain words onto other words and sentences onto other sentences. Nothing bars deflationists—or any other theorist of truth—from using such a mapping. Differences emerge over the *philosophical* questions of whether—and how—one explains core concepts employed by the machinery—namely, reference and truth. Otherwise put, the philosophical question is whether one thinks *the machinery must be grounded*—and if so, how it is grounded. Deflationists claim that reference and truth are not in need of substantive metaphysical explanations. This means, in effect, that they think the machinery of standard semantic theory needs no grounds. Given the empirical success of semantic theory, this seems—at least to me—a highly implausible stance. For where a theoretical apparatus leads to successful explanation and prediction, there is typically a reason for

¹⁷ For discussion of complications here see D. Bar-On, C. Horsick, and W. Lycan, “Deflationism, Meaning and Truth-Conditions”, reprinted in B. Armour-Garb and Jc Beall (eds.) *Deflationary Truth*, (Chicago, IL.: Open Court, 2005); and D. Patterson, “Deflationism and the Truth-Conditional Theory of Meaning”, *Philosophical Studies*, 124 (2005) 271–94.

why this is so. There is something to say about what the apparatus is tracking.

Representationalists and antirepresentationalists about language (and mind) on the other hand, disagree over what that ground is: they disagree over the nature of meaning and reference. Both can agree that meaning consists in truth-conditions and employ the machinery of formal semantics. But one believes this must be cashed out in terms of a correspondence theory of, while the other explains them in terms of verifiability or superwarrant.

Given these facts, it is not hard to see that alethic functionalism allows us to employ the machinery of formal truth-conditional semantics as well. But the fact that truth is (at least possibly) variably manifested also allows us to hold that meaning is variably manifested, and moreover, appeal to the respective manifesting properties to *explain what grounds the formal machinery of semantic theory*. That is, as we've already seen in our discussion of content determination, alethic functionalism can allow that for some types of sentences, their meaning is ultimately determined by the conditions under which they are superwarranted. For other types of sentences, their meaning is given by the conditions under which they causally isomorphic to map various objects and properties in the world. The functionalist theory does not hinge on either of these philosophical accounts of the nature of meaning being successful, but is compatible with the thought that, as we might put it, each such project *explains how meaning as such—truth-conditions—is manifested* for some types of sentences.

5. Superwarrant and Reference

Above I argued that one could hold that content (and meaning) are at least partly determined by truth-conditions, but also hold that truth is manifested by superwarrant. Moreover I also noted that REPRESENT itself can be read as a truism or platitude. These three ideas might be thought to be in immediate tension. For

where truth is superwarrant, we will definitely not understand the truth of a belief as hinging on whether its component parts denote certain objects. Denying that truth consists in such a mind/world relation, after all, is part of the point of accepting a superwarrant theory of truth.

In point of fact, however, antirepresentationalist theories of truth can accept REPRESENT and similar truisms linking truth with reference or denotation. Such truisms may include, for example, where $\langle c \rangle$ is a particular concept expressed by a singular term:

$\langle n \rangle$ denotes if and only if $n = x$.

Likewise for concepts expressed by predicates:

$\langle F \rangle$ is true of x if and only if x is an F .

Such truisms will presumably be accepted by any philosopher, but how they are understood and explained will vary. As noted in the last section, deflationists will take them (or their instances) as basic, in need of no explanation. Traditional monists will naturally look for an underlying theory of denotation that explains why the truisms are true. Semantic functionalists will take it that they are part of a functional theory of the denotation relation, a relation that will be open to multiple manifestations, depending on the domain of concepts in question. For the functionalist, the deflationist is only partly right: they are right that there need be no single metaphysical explanation for why all given concepts denote. But they are wrong to think there is no explanation at all. Rather, what explains the denotation of a given concept depends on the type of concept it is, and the way that denotation is manifested by concepts of that type.¹⁸

¹⁸ I have here restricted my remarks to conceptual reference. Related notions like satisfaction will likewise be treated similarly by the functionalist. At least, I see nothing to bar the functionalist from holding that what makes it the case that a given predicate is satisfied by a sequence of objects depends on the type of concept that predicate expresses. That is, just as reference is—if the functionalist is right—open to a functionalist analysis, so too is satisfaction.

How will that relation manifest itself differently? A full treatment of this question, obviously, merits a book in itself. But the following overall remarks seem plausible. An essential difference between domains where truth is manifested by superwarrant and domains where it is manifested by something like correspondence is in the *order of explanation* with regard to denotation/reference and truth. Representationalists who adopt accounts like CAUSAL will take it that:

The belief that x is F is true *in virtue* of the object denoted by our concept of x having the property denoted by our concept of F.

Here the explanation goes from being to denotation to truth, as it were. But someone who is attracted to the idea that in some domains, superwarrant manifests truth will hold that in domains where it does so,

The object denoted by our concept of x has the property denoted by our concept of F *in virtue* of the truth of the belief that x is F.

Where this in turn amounts to:

The object denoted by our concept of x has the property denoted by our concept of F in virtue of the belief that x is F being superwarranted.

In short, in domains where truth is manifested by superwarrant, denotation is explained by truth. No doubt, such a view is not plausible across the board. It is not plausible, for example, that trees have leaves just in virtue of the fact that my belief that they do is warranted and will remain so under all increases of information. But it is far from clear we must think that this holds for all kinds of beliefs. And even if we do think it holds for all kinds of beliefs (because, for example, we think that beliefs are by definitional representational) then it is far from

clear that it holds for all truth-apt judgments. In any event, the present point is that in domains where superwarrant manifests truth, content determination and related notions like reference or denotation will themselves be explained in terms of the conditions under which beliefs are superwarranted. This allows us to retain the thought that content is determined by truth-conditions and yet also retain the insights behind verificationist views of content.

The situation is further complicated by how one understands “warrant” in superwarrant. In Chapter 2, we noted that one plausible way of understanding warrant is in terms of what we called supercoherence:

Supercoherence: a belief is supercoherent if and only if that belief is a member of a supercoherent set of beliefs, where a system is supercoherent when it is coherent at some stage of inquiry and remains so without defeat in every successive stage of inquiry.

Where such an account of warrant is merited, truth is superwarrant, and TCD is accepted, content is determined by a belief’s relationships to other beliefs. Such an account of content is therefore committed to some form of holism. Yet the advocate of this view of belief content in some domain can still accept REPRESENT. Their view requires a euthyphronic reading of that biconditional, one that claims that concepts have their denotations in virtue of the beliefs that contain those concepts being member of a (super)coherent set of beliefs. In effect, this last sort of position introduces the basic principles of conceptual role semantics under a truth-theoretic guise. For it would be natural on such a position to hold that concepts have their content in virtue of how they contribute to the beliefs that they compose (a version of Frege’s context principle)—that is, in virtue of their role in inference, their conceptual role in short.

Whether any of the above position on denotation is ultimately plausible is not what is at issue here. The point is that nothing bars

a functionalist about truth from also holding that the concept of denotation is itself functionally analyzable in terms of certain truisms. Moreover, nothing bars the philosopher who thinks that truth is, in some domains, manifested by superwarrant, from holding—and explaining—those truisms. In domains where truth is manifested by an epistemic property like superwarrant, representation and denotation are themselves to be understood in terms of superwarrant—and therefore truth. In domains where truth is manifested by correspondence, the order of explanation is reversed: in those domains, whether a given belief is true depends on its referential properties.

6. Belief and Truth-Aptness

Not everything we say or think is true or false; indeed not everything we say or think is even a candidate for truth or falsity, or truth-apt. Commands aren't true or false, neither are questions, or cheering and booing. On this everyone agrees. From here, discussions about truth-aptness tend to be framed by two extremes: those who think it is easy for something to be truth-apt and those that believe it is, in a variety of ways, a more difficult achievement.

Where the alethic functionalist stands on this issue depends on what sorts of things one takes to be truth-apt. It is maximally easy on my view for propositions to be truth-apt, since propositions are just the sort of thing that is true and false. Discussions over truth-aptness, however, are typically concerned not with propositions per se, but with the conditions under which a proposition is expressed by an assertion or declarative sentence. Those like Crispin Wright who believe it is easy for a sentence to be truth-apt—sometimes called “minimalists”—reason that if it walks like a duck, and quacks like a duck, it's a duck. That is, if a given sentence, for example, can be negated, used to form the antecedent of a conditional, and meets other standards of syntactic discipline as it is sometimes put, then it meets the very minimal standards for

being truth-apt.¹⁹ There is nothing more to it. Other philosophers think that it is not so simple: some sentences (and utterances of sentences) that may seem to be truth-apt are not because they fail to meet some further condition. Thus, some might hold, for example, that

- (a) If an assertion expresses a proposition, it conveys a representational mental state.
- (b) Ethical assertions (e.g. “murder is wrong”) do not convey such states.
- (c) Therefore, ethical assertions do not express propositions, and hence are neither true nor false.

Arguments of this basic sort have classically been used to argue for expressivism about ethics. Something like claim (b) is sometimes thought to follow from so-called Humean views of motivation (we’ll have occasion to examine this point in the next chapter). Claim (a) is sometimes thought to be a platitude, as I noted myself in Chapter 2; together, they post a restriction on what sorts of assertions can be truth-apt.

The picture is more complicated by “in-between” positions like Simon Blackburn’s quasi-realism. Blackburn rejects (a) but holds (b). That is, Blackburn is perfectly willing to grant that my ethical assertions can express propositions and be truth-apt. According to Blackburn, this is because truth is essentially deflationary in character. As he says, “we can talk of metaphysical, mathematical, modal truth because that is just to repeat our commitments in these areas. But what is going on when we have such a commitment is to be understood in other terms”.²⁰

Blackburn, like Mark Timmons, another prominent ethical expressivist, is willing to grant that moral commitments convey beliefs, in a minimal sense of “belief”.²¹ A natural way to understand this suggestion, in the present context, is to take it as implying that

¹⁹ See his *Truth and Objectivity* (Cambridge, MA.: Harvard University Press, 1992).

²⁰ S. Blackburn, *Ruling Passions* (Oxford: Oxford University Press, 1998), 78.

²¹ Blackburn *Ruling Passions*, 79–80.

beliefhood is subject to a functionalist analysis of the sort provided by truth in this book. According to such an account the concept of belief is explicated functionally in terms of its relations to other concepts via the now familiar truism strategy. In the case of belief, these will include, most relevantly,

To believe that *p* is to hold the proposition that *p* to be true.

A sincere assertion of a sentence like “snow is white” expresses the belief that snow is white.

Other platitudes, familiar from functionalist accounts of belief in the philosophy of mind, will connect beliefs with sensations, other mental states (like desires) and action.

Details aside, the present point is that the functionalist account of belief does not require that beliefs be *representations*. What serves to distinguish beliefs from desires or hopes, on this view, is not that beliefs represent. It is that, among other things, to believe is to hold as true. Some of the mental states that are beliefs may indeed be representational. But on this view, nothing in the concept of belief requires that to be the case.

On this point, functionalism is consistent with Blackburn’s view, according to which moral beliefs are not representations; they don’t map the world. For Blackburn, representing the world in a certain way is one sort of commitment, but it is not the only one. I can also be committed in sentiment, as it were. Unlike representations sentimental commitments project our attitudes upon it. Consequently, the propositions to which I am ethically committed can be true and false, but they do not represent the world like those propositions I can believe.

The principle disadvantage of quasi-realism is that it threatens to be unstable. The source of the instability is the endorsement of deflationism on the one hand and the wish to accommodate semantic diversity on the other. The quasi-realist appeals to deflationism to accommodate the fact that ethical assertions, for example, certainly appear to be truth-apt. They figure as antecedents in conditionals and can be inferred in valid arguments.

Thus the source of semantic diversity between the moral domain and others, for the quasi-realist, lies with the differences in the types of commitments moral assertions convey. Ethical assertions don't convey representations but sentiments. But if moral sentiments/beliefs are not representations, then one must have a conception of representational belief that distinguishes it from sentiments. The most natural thought is that a representational belief is different from, say, a desire, in that it involves a particular direction of fit. Representational beliefs map the world. But representations can be correct or incorrect, accurate or not. And here we find it overwhelmingly tempting to say that an accurate representation is a true representation. But if we do say this, and it seems very plausible, then we have appealed to a non-deflationary conception of truth to explain the concept of representational belief. Hence the instability: the quasi-realist appeals to deflationism to explain the cognitive unity of ethical assertions with other forms of assertions. But it seems likely that one can only make heavy-weather of the difference between representational beliefs and other propositional attitudes by appealing to a non-deflationary account of truth.

It is tempting to put it this way: according to the quasi-realist, one can use the *word* "true" with any sort of assertion, but only assertions which express representational beliefs can really be true.

Functionalists, like quasi-realists, take the various domains that Blackburn mentions to be both semantically diverse and yet cognitively unified. But functionalism avoids the disadvantages of quasi-realism while leaving room for the intuitions that motivate the view in the first place. The difference between the two approaches hangs on how each view understands both the diversity and the unification. The quasi-realist sees the unification as emerging out of the fact that truth is a mere logical device for semantic ascent and generalization. To claim that a commitment is true is to simply repeat that commitment as Blackburn says. The matter is more complicated for the functionalist, who requires that the commitments of the domain in question supports the existence of a property that can play a more complex role. Ironically, it is therefore much more of an open

question whether some domain of commitments which may *appear* truth-apt really is truth-apt. In this sense, the functionalist leaves more room for non-cognitivism than the quasi-realist—so long as she is willing to grant (a) and (b) above, she can certainly grant (c).

But as the functionalist allows for the possibility that there can be propositions that we can believe in a non-representational way but that are nonetheless truth-apt. *Nothing in her view precludes her, with Blackburn, from denying (a) but holding (b).* Moreover she can say this without the threat of instability that endangers quasi-realism. For on her view, ethical propositions could be true but have their truth realized by a very different sort of property that realizes the truth of propositions that are the objects of our representational beliefs. One is free to hold that some beliefs (and believed propositions) represent or map the world and are true insofar as they do so accurately, or correspond to its natural geography. But in so doing, one does not invoke a distinct theory of truth, *but rather a particular theory of what manifests truth.* Moreover, there is no sense that the ethical commitments, should they have their truth manifested in a distinct way, are somehow less than fully truth-apt. They are truth-apt just insofar as they are capable of having a property that plays the truth-role. If, as I show in the last chapter of this book, it can be maintained that a property other than correspondence does so in the ethical domain then they are as true as any other type of proposition.

7. How does Pluralism Account for the Normativity of Truth?

We've noted that two truisms about truth are these:

Truth is a worthy goal of inquiry.

That is, it is *prima facie* good that we indirectly aim at believing what is true and (arguably) only what is true when pursuing answers to the various questions that interest us.

Truth is the norm belief.

That is, it is correct to believe a proposition only if it is true. Truth is the “good” in the way of belief, as James put it.

Both of these truisms state normative or evaluative facts about truth. In earlier chapters, we noted that both claims require that the value involved is general in character. Moreover, we noted that it is highly plausible that truth is not just an instrumental goal of inquiry, but its ultimate goal or end.

The functionalist theory accommodates these demands. First, the normative generality of both truisms is accommodated by the immanent functional property being univocally applicable to any proposition, no matter what its domain. Thus we may say without equivocation that, other things being equal, a proposition is correct to believe when it is true. According to functionalism, this means that it is correct to believe propositions that manifest truth. Likewise, we can say that no matter what sort of question we are asking, what sort of inquiry we are engaging in, it is a good to believe propositions that are true—or have the immanent functional property of truth.

And just as clearly, the functionalist can hold that believing propositions that have that property is a cognitive end in itself and not, for example, a means to some other cognitive end.

According to functionalism, both normative truisms about truth are integral to what truth is. They are part of the core folk-theory of truth that individuates the truth-role. Consequently, any property that manifests truth must satisfy these normative platitudes. So for example, any property that plays the truth-role for propositions of a particular domain must be such that it is correct to believe propositions that have that property. Crucially, however, *this needn't be because of any intrinsic normative facts about the manifesting property itself*. Such properties considered independently of their role in manifesting truth, may be fully “descriptive”. That is, correspondence qua correspondence may have no normative features. It may only be that correspondence qua manifestation of truth has such features. In other words, what makes it correct to believe propositions with the relevant manifesting property

is precisely that it manifests *truth*—something that a number of different properties might do. A proposition is true when it has a property that manifests truth. Properties manifest truth when they play the truth-role, and part of playing the truth-role is being a property that makes propositions of a particular domain correct to believe. Thus it is correct to believe what manifests truth in some way *just because in doing so, one believes what is true*.²²

Of course, none of this explains *why* truth is a worthy goal of inquiry and why it is a Norm of Belief. Neither does it explain what either value consists in. But that is not surprising, since those are other matters altogether.²³

8. Do we need Pluralism about Truth—Why not a Simple Pluralism about Content?

A *prima facie* appeal of functionalism is that it allows us to make sense of the idea that truth is many, by making sense of the idea that there can be more than one property that makes propositions true. This in turn allows us to accommodate both the semantic diversity and the cognitive unification of thought. One might think that one could do this much more easily by endorsing a pluralist view of *content*. The thought would be to explain semantic diversity by appealing not to different kinds of truth but to different kinds of *truths* or different kinds of propositions.

As the above discussion of quasi-realism indicates, this is easier said than done. For one must explain the difference in which the different kinds of propositions consist in a way that does not appeal—even in part—to the different ways in which propositions can be true. This is difficult to do while at the same time respecting cognitive unity.

²² Compare: What makes it good to avoid being in some state that realizes the pain-role is that by doing so, one avoids being in pain. In both cases, what matters—the normative force—depends on facts about the role.

²³ For answers to these and other questions about the normativity of truth, see M. P. Lynch, *True to Life* (Cambridge, Mass.: MIT Press, 2004).

Consider the two most natural ways one might try to explain semantic diversity by appeal to distinct kinds of content while retaining cognitive unity. The first method achieves cognitive unity on the cheap by appealing to a deflationary view of truth. All types of propositional content are truth-apt in the thin minimal sense of ‘true’ endorsed by deflationism. Semantic diversity is then explained by appeal to allegedly truth-independent differences in that content. But what differences? As I just indicated, I doubt that one can hope to explicate those differences by appeal to differences in the psychological attitudes we hold towards the relevant kinds of propositions. For one will find oneself appealing through the backdoor to “thicker” notions of truth and representation in order to explain the differences in attitude. Neither will it do to simply say that ethical propositions, for example, don’t “describe” the world while propositions about the physical world do. For, again, how does one explain what makes one proposition a description without appealing to the fact that descriptions are correct when they correspond to facts? If on the other hand, one insists that correct descriptions are simply true descriptions in the deflationary sense of “true” then the motivation for denying the status of “description” to ethical propositions has been summarily removed.²⁴

The second method looks not to deflationism to supply cognitive unity but a thicker theory of truth, such as the correspondence account. Semantic diversity is then explained indirectly by the thought that “different sorts of propositions correspond to different sorts of facts”. Ethical propositions correspond to ethical facts; mathematical propositions correspond to mathematical facts. For each type of proposition there is its own brand of fact.²⁵

²⁴ Mark Timmons proposes a hybrid view: moral assertions do not describe the world but are nonetheless capable of being true in the deflationary sense. Yet Timmons, unlike Blackburn, appeals directly to a distinct notion of truth as correspondence to explain the difference between descriptive and non-descriptive claims. Consequently, Timmons view ends up being closer to a type of pluralism about truth. See his *Morality without Foundations* (Oxford: Oxford University Press, 1998).

²⁵ See G. Sher in “Functional Pluralism”, *Philosophical Books*, 46 (2005), 311–30, and “In Search of a Substantive Theory of Truth”, *The Journal of Philosophy*, 101 (2004), 5–36.

Two difficulties immediately present themselves. The first is that it is not clear whether the proponent of such a view is explaining semantic diversity so much as explaining it away. But put this aside: the real problem with any such position, as I see it, is that it is difficult to see how any robust correspondence theory of truth—as opposed to a simple endorsement of the *Objectivity truism*—can hope to solve the scope problem. As I argued in Chapter 2, robust notions of “fact” and “correspondence” have a way of favoring some kind of claims over others. Understand correspondence causally, for example, and it is unclear how propositions about numbers and norms can be true.²⁶ Consequently, it is unclear how the view can avoid the usual fate of those who think that all sorts of assertions are apt for correspondence and only correspondence truth: adopting an error theory towards those domains that stubbornly refuse to fit any facts in the endorsed robust sense of fact. Cognitive unity—the thought that our different sorts of thoughts can all be true—threatens to slip away.

Nothing has been said here that *proves* that neither of these methods for respecting cognitive unity and semantic diversity must fail. But neither appears particularly promising: it looks likely that both approaches will end up abandoning one intuition or the other. Hence to the extent we wish to respect both intuitions, to that extent a view like functionalism—a view that allows for alethic pluralism—is the better view.

9. A Circular Analysis?

The remarks in this chapter have been deliberately programmatic and suggestive; I have not attempted to develop semantic

²⁶ One possibility here is for the advocated of correspondence to advocate a contextualist view like Horgan’s: see in T. Horgan, “Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence”, in M. Lynch (ed.), *The Nature of Truth* (Cambridge, MA.: MIT Press, 2001), 67–96. See also T. Horgan and R. Barnard, “Truth as Mediated Correspondence”, *The Monist*, 89 (2006), 28–49. For criticisms of this approach, see my “The Truth in Contextual Semantics”, *Grazer Philosophische Studien*, 63 (2002), 173–95. Horgan responds in the same volume.

functionalism so much as to describe it in rough outline. Nonetheless even at this level of generality, one might think that there is an obvious objection to the account:

You've said that alethic functionalism can explain, e.g. content in terms of truth. Yet the theory seems to require that we have a prior grasp of content—and other concepts—in order for us to determine which property manifests truth in a given domain. You can't have it both ways: either content, reference and so on explains truth or truth explains those other concepts.

It is the case that alethic functionalism demands that we have a prior grasp on, e.g. what proposition we believe, before being able to determine which property manifests truth for what we believe. But this is no more troublesome than the obvious fact that I must have a prior grasp on what you believe in order to assess whether it is true. In neither situation must we presuppose a *philosophical account* of content (or content determination) in general, which is what would have to be the case were the alleged circularity to have any bite.

Nonetheless, the objection does usefully point to a feature of the approach being advocated here. My view is that truth is a kind of functional property. It is defined by certain truisms that relate it to other properties such as proposition, belief—and content. Insofar as some of our truisms about truth relate truth to these other properties, we use these properties to help us understand what truth is. If, as I've been suggesting in this chapter, we go still farther and embrace semantic functionalism, we can then turn around and use truth to explain these same properties. Here as elsewhere, light dawns gradually over the whole.²⁷

²⁷ L. Wittgenstein, *On Certainty* (Oxford: Blackwell, 1969), sect. 141.

8

Applying the View: Moral Truth

... The question is not whether [“true” and “false”] are in practice applied to ethical statements, but whether, if they are so applied, the point of doing so would be the same as the point of applying them to statements of other kinds, and if not, in what ways it would be different.

Michael Dummett, “Truth”

I. Back to the Puzzle

I began this book with a puzzle. We judge and believe many sorts of propositions to be true, from those about the ordinary objects cluttering our daily life to ones about numbers, justice, and beauty. Yet, if truth is always and everywhere the same—if, in particular, it is everywhere a matter of correspondence with reality—it is puzzling how all these different sorts of judgments can be true. It seems that with regard to some types of judgments, judgments of that type can be true, and yet we can’t explain this fact. Hence the puzzle.

Nowhere is the puzzle more gripping than in the case of morality. This is because the conditions that make the puzzle puzzling are particularly salient with respect to moral judgments. Moral propositions seem truth-apt and yet it is mysterious how this can be so.

There are at least three widely acknowledged reasons for thinking that moral judgments are truth-apt. None of these reasons is completely persuasive absent further considerations. But they are persuasive enough to make a *prima facie* case—as indicated by the fact that those who deny that moral judgments are truth-apt typically expend considerable effort trying to explain these reasons away.

The first *prima facie* reason to think that moral judgments are truth-apt is that, like any other judgment, they have cognitive surface-structure. A judgment has cognitive surface-structure when it can be meaningfully negated, when it can figure in detachable conditionals, and generally be understood, for purposes of logic, as truth-functional. Moral judgments have all of these features.

A second reason to think that moral judgments are truth-apt is that they are subjected to norms of epistemic appraisal. Moral judgment takes place in what Sellars called the “space of reasons”. If I make a moral judgment about, e.g. the injustice of the death penalty, and you deny that judgment, in the normal course of things, I am committed to giving you a reason or some other evidence in favor of my judgment. If I cannot produce a reason or any other evidence—even a reason that is itself up for challenge—I should retract my judgment, or at least lower my confidence in it. This makes most moral judgments different from many judgments of taste. I need not retract my judgment that fettucine alfredo is delicious if you deny it. Nor am I committed to giving you a reason for why I made the judgment. But if I declare that it is sometimes right to torture prisoners, I am committed to giving you a reason if you deny that judgment. If I cannot produce one, I incur an obligation to, at the very least, lower my confidence in that opinion.

The third reason moral judgments seem truth-apt is that they have objective pretensions. That is, we normally take ourselves to be capable of making moral mistakes and being in moral ignorance. Indeed, we typically think of moral growth and maturity as a process that involves identifying and correcting past moral error and coming to appreciate morally relevant factors of which one had previously been unaware. Thus someone who was raised in

a racist environment might later come to see his previous views as morally mistaken. In doing so, he seems to commit himself to thinking that his previous views were not only false, but that their truth-value did not depend on him judging them to have a value.

All three of these reasons can be challenged; but together they represent at least a *prima facie* case for taking moral judgments to be truth-apt. The question is just what sort of truth they are apt for. Our puzzle becomes particularly gripping if we assume they are apt for correspondence truth. This would be to take it that if our moral judgments are true, then they are so by virtue of corresponding, in some sense, to mind-independent moral objects and moral properties.

The problem becomes acute if “correspondence” is understood in terms of representation. For it is unclear how moral judgments *could* represent objects and properties in the world if we understand representation in either of the two substantive ways explored in Chapter 2. There I noted that if judgments are to correspond, they must respond. That is, representational theories of truth are only plausible with regard to the G-ish domain where we show that

Responsiveness: Mental states with G-ish content are causally responsive to an external environment that contains Gs.

That is, you can't map what isn't there and you can't map well that with which you don't have some even indirect causal contact. But there are well-known and intuitive reasons for thinking moral properties, should they exist, would not be the sort of properties with which we enter into causal contact.¹ The problem, as Crispin Wright has emphasized, is not just that one might wonder how moral properties can be thought to causally explain our moral beliefs. The problem is that it is hard to see how moral properties

¹ See e.g. G. Harman, *The Nature of Morality* (Oxford: Oxford University Press, 1977); S. Blackburn, “Just Causes”, *Philosophical Studies*, 61 (1991), 11–15. For responses, see P. Bloomfield, *Moral Reality* (Oxford: Oxford University Press, 2001); and D. Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), 182–97.

can be seen to be part of a causal explanation for *anything other* than moral beliefs.² Wherever we are confident that we are cognitively responsive to physical objects, we typically take it that the objects to which we are so responsive cause more than just our responses. We respond to the presence of cats on mats by forming beliefs about cats, but the cat's being on the mat causes more than my belief—it also causes the mouse to stay in the hole, and a flea to be on the mat, and so on. This nexus of causal connections is part of what convinces us that there is something to which we are responding when we form our beliefs about cats. But with Wright and Harman, one might wonder whether this wide spectrum of causal connections exists in the moral case.³

The point here isn't that these considerations should make us give up on the existence of moral properties, although they are precisely the sort of considerations that have led many to do so. Rather, my point is that such properties, if they exist, *seem unsuited to being represented according to our best theories of what representation is*. If so, and if, as I argued earlier, "correspondence" is best understood in terms of representation, then we have reason to doubt whether moral judgments are true way of correspondence.

Of course, one might reply that while moral judgments are true in virtue of corresponding to moral properties, the type of correspondence involved is not a matter of accurate representation. This is not unreasonable, but it requires a theory of correspondence that is both non-vacuous and appropriate for moral judgments. As I argued in Chapter 2, avoiding vacuity is a major challenge for correspondence theories. Saying that, e.g., a moral judgment

² See C. Wright, "Truth in Ethics" in his *Saving the Differences* (Cambridge, MA.: Harvard University Press, 2003), 199.

³ The so-called "Cornell Realists" can be seen as challenging this assertion. If moral properties just reduce to natural properties, then perhaps our beliefs might causally map them after all. See N. Sturgeon, "Moral Explanations", in Copp and Zimmerman (eds.), *Morality, Reason and Truth* (New Jersey: Rowman & Littlefield, 1984); R. Boyd, "How to be a Moral Realist", in G. Sayre-McCord (ed.), *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988). A well-known criticism of such views with which I am sympathetic is T. Horgan and M. Timmons's "moral twin-earth" argument; "New Wave Moral Realism Meets Moral Twin Earth", *Journal of Philosophical Research*, 16 (1991), 447–65.

is true if and only if it corresponds to the moral facts is simply to repeat the Objectivity truism in another way. It doesn't amount to a substantive theory of truth.

The above considerations hopefully illustrate why the puzzle with which we began is particularly pressing in the case of morality. We have reasons to believe that moral judgments are truth-apt. But it is difficult to explain how this can be so. It is particularly difficult if we think that moral judgments are made true by correspondence. Hence we face a choice. We can give up the truth-aptness of moral judgments—despite our initial reasons. This is the route of the traditional expressivist.⁴ We can accept that moral judgments are true or false by virtue of correspondence, but declare them all false. This is the route of the traditional error theorist.⁵ We can, perhaps, develop a non-vacuous theory of correspondence truth that makes sense of how moral judgments correspond. As I just noted, this seems difficult.

Another option we could try is going deflationist across the board. This would be to claim, in effect, that nothing makes moral judgments true because nothing makes judgments true. There is nothing substantive to say about how moral judgments are true. Needless to say, this route is hard to combine with a robust realism about morality. For if there are moral objects and properties out there in the world, it would be curious indeed if they had nothing to do with why the judgments we make about them are true. Not surprisingly then, those inclined to go deflationary are more likely to endorse a form of antirealism, such as Blackburn's expressivist quasi-realism.⁶ But this too has serious costs. First, it has all the costs of deflationism itself. Among other things, committing oneself to deflationism means committing oneself to the dubious view that

⁴ Classic presentations include, A. J. Ayer, *Language, Truth and Logic* (New York: Dover, 1952); and C. L. Stevenson, "The Emotive Meaning of Ethical Terms", *Mind*, 46 (1937), 14–31.

⁵ J. L. Mackie, *Ethics: Inventing Right and Wrong* (New York: Penguin, 1991).

⁶ S. Blackburn, *Ruling Passions* (Oxford: Oxford University Press, 1998); see also Timmons *Morality without Foundations* (Oxford: Oxford University Press, 1999). As noted in the last chapter, however, Timmons might be better construed as a pluralist.

truth isn't useful for explaining things like content or meaning. Second, and as I argued in the last chapter, any combination of such deflationism with moral antirealism—some version of other of expressivism say—is inherently unstable.

In any event, I want to put those options aside here, and instead take up the following question. If moral judgments were true in a non-deflationary sense, but not by way of correspondence, what property *would* make them true—that is, what property manifests truth for moral judgment? In other words, if we assume that correspondence and deflationism are off the board as accounts of moral truth, what options remain? This amounts to exploring whether there is another solution available to our puzzle about moral judgments. Addressing it allows us to not only test the functionalist theory, it helps to illustrate how adopting that theory can help us to contribute to long-standing philosophical debates.

2. The Nature of Moral Warrant

Where we despair of understanding truth in terms of representation, the obvious alternative is some form or other of what I called in Chapter 2 antirepresentationalism. This route seems particularly plausible here.⁷ If, as we are assuming, moral judgment takes place within the space of reasons—we provide reason and evidence for our moral beliefs—our moral judgments are subject to significant rational norms. Consequently, a natural suggestion is that moral judgments are made true by a property that is constructed out of those epistemic norms.

A first step to understanding the character of such a property is to examine the nature of moral warrant. In explaining his theory of justice, John Rawls influentially suggested that:

Here the test is that of general and wide reflective equilibrium, that is, how well the view as a whole meshes with and articulates our more firm

⁷ In the appendix to this chapter I offer some further considerations for thinking that what manifests moral truth should be antirepresentational in character.

considered convictions, at all levels of generality, after due examination, once all adjustments and revisions that seem compelling have been made. A doctrine that meets this criterion is the doctrine that, so far as we can now ascertain, is the most reasonable for us.⁸

Rawls' concern was theory-acceptance, but subsequent thinkers have broadened his insight and cast it as a theory of justification in ethics generally.⁹ The basic thought is that a moral judgment is warranted to the degree that it is woven tightly into the rest of the moral fabric, to the degree, to speak more plainly, that it *coheres* with our considered moral judgments and relevant non-moral convictions. Rawls suggestion, in other words, is often taken to underwrite a broadly coherentist approach to moral epistemology.

The basic approach has won wide support amongst realists and antirealists alike. Thus, we find David Brink—an arch-realist and stout defender of objective moral facts—opting for moral coherentism over foundationalism:

We all have or entertain moral beliefs of various levels of generality ... many of these moral beliefs depend on other moral beliefs. For instance, beliefs about the value of a particular activity depend, among other things, on ideals of the person (i.e., moral beliefs about what kind of persons we ought to be). Moral beliefs also depend on nonmoral beliefs. For instance, beliefs about the moral or political legitimacy of a welfare state depend on nonmoral beliefs about such things as human nature, social theory, and economics. A coherence theory of justification in ethics demands that these and other beliefs be made into a maximally coherent system of beliefs.¹⁰

Coherentism's principle rival is foundationalism. According to foundationalism broadly construed, judgments come in two kinds:

⁸ J. Rawls, "Kantian Constructivism in Moral Theory", *Journal of Philosophy*, 77 (1980), 534.

⁹ The tip of the iceberg would be: N. Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics", *Journal of Philosophy*, 76, 256–82; S. Scheffler, "Justification and Commitment", *Journal of Philosophy*, 51, 180–90; G. Sayre-McCord, "Coherentist Epistemology and Moral Theory", in W. Sinnott-Armstrong and M. Timmons (ed.), *Moral Knowledge?* (Oxford: Oxford University Press, 1996); and Brink, *Moral Realism and the Foundations of Ethics*.

¹⁰ Brink, *Moral Realism*, 102.

(a) those warranted by other judgments; and (b) those basic or foundational judgments warranted in some other way. Thus for empirical judgments, a foundationalist typically takes it that there are some judgments that are warranted just in virtue of being the product of reliable perceptual mechanisms. Coherentism can be minimally defined as the rejection of (b), and hence an endorsement of the claim that all judgments must be warranted by other judgments. In this sense, all warrant is, for the coherentist, “inferential”.¹¹

The contrast with empirical judgment suggests why realists and antirealists about morality alike take coherentism as a plausible moral epistemology. A moral foundationalist must maintain that some judgments are warranted by something other than another judgment. Thus either some moral judgments are simply self-evident or some moral judgments are justified by something analogous to perception—moral “intuition”. Given the seemingly inherent “contestability”—to use Wright’s phrase—of many of our moral commitments, neither approach has garnered many followers.¹²

So coherentism is an attractive and entrenched moral epistemology. According to the theory, a moral judgment is warranted to the degree that the framework to which it belongs is coherent. Such systems, as Brink emphasizes, include both moral and relevant non-moral judgments. Call such systems *moral frameworks*. But what does it mean to say a moral framework is coherent?

Coherence theorists frequently apologize for the fact that coherence is notoriously resistant to precise characterization. In fact, no apology is necessary. Its resistance is the unsurprising result of the fact that “coherence” names a family of epistemic desiderata. That

¹¹ Defenses of coherentism as a general theory of justification can be found in Bonjour, *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard University Press, 1985); G. Harman, *Change in View* (Cambridge, MA: MIT Press, 1986); W. Lycan, “Plantinga and Coherentisms”, in Jonathan Kvanvig (ed.), *Warrant in Contemporary Epistemology* (Totowa, NJ: Rowman and Littlefield, 1996).

¹² Notable recent exceptions include: Shafer-Landau, *Moral Realism: A Defence* (Oxford: Oxford University Press, 2003), ch. 11 and R. Audi, *The Good and the Right* (Princeton, NJ: Princeton University Press, 2005).

is, a framework is coherent insofar as, and to the degree to which, it exemplifies the following virtues:

- Mutual explanatory support: A framework's judgments are mutually explanatory when (a) they are explanatorily compatible, none acts as a defeater for another; and (b) they are explanatorily connected: each judgment within the framework is positively supported, either inductively, abductively, or deductively by the others in the framework.
- Predictive power: the framework is a reliable predictor of future experience and judgment.
- Simplicity: the framework's explanations are neither ad-hoc nor needlessly complex.
- Completeness: the framework contains, for every proposition of the relevant kind, either a judgment in that proposition, or a judgment in its negation.
- Consistency: judgments within the framework are not logically inconsistent.

Call these *coherence-making features*. Such features themselves come in degrees: members of a framework can be more or less consistent,¹³ more or less mutually explanatory, etc. A framework of judgments increases in coherence to the degree to which it exemplifies these features, on balance, to a greater degree. "On balance" because the features are not themselves isolated in their coherence-increasing power. A framework would not be more coherent on balance, for example, simply by increasing its size (completeness) by including consistent but explanatorily unconnected judgments. Intuitively, by increasing its explanatorily isolated judgments, the coherence of the framework would on balance remain static or decrease. On the other hand, a framework would be *maximally coherent*, presumably, were it, on balance, to exemplify as many of these features as possible to the greatest degree possible, where

¹³ A system is more or less consistent to the degree it which it contains inconsistent pairs of judgments.

the limit of possibility here is determined by, among other things, the nature of judgments in question.¹⁴ Thus it may be that for some types of judgments, full explanatory connectedness is not possible. If so, then some framework may qualify as maximally coherent even if they are not positively supported by other judgments in the framework (although no other judgment acts as a defeater for them either). Likewise, for some domains—such as the moral domain—predictive power may have less weight, and completeness may be impossible. More on this last point shortly.

Given coherentism's plausibility as an account of the structure of warrant in morality it seems natural for anyone attracted to an epistemic account of the property that makes moral judgments true to appeal to it. Moreover, a marriage between coherentism about warrant and coherentism about truth is happy for another reason as well: it would answer one of the most pressing objections to coherentism in epistemology. It is a platitude that judgments are warranted only if they are likely to be true. Why should the fact that a judgment is a member of a coherent framework of judgments make it likely to be true? If we avail ourselves of a coherentist theory of what makes moral judgments true, we can answer: coherent moral judgments are likely to be true because what makes moral judgments true is itself constructed out of coherence.

3. From Coherence to Supercoherence

In the last section I suggested that a moral judgment is warranted to the degree that the framework to which it belongs is coherent. How might we use this notion or warrant in the moral realm to construct a theory of truth?

¹⁴ Thus none of the desiderata are strictly speaking necessary. The present account is therefore open to the possibility that some maximally coherent systems of judgments might contain some contradictions. See G. Priest, *Doubt Truth to be a Liar* (Oxford: Oxford University Press, 2006), 49–51.

A natural, if wrong-headed, suggestion is that true moral judgments belong to *maximally* coherent moral frameworks.¹⁵ It is natural because it would distinguish between mere warrant and truth, while still constructing the latter out of our standards for the former. Mere warrant is a matter of degree; truth is not. And warrant is an achievement which is a step along the path of the greater achievement of truth. Our initial suggestion allows for both points: belonging to a maximally coherent framework is an absolute property; and maximal coherence is an achievement greater than coherence.

But the suggestion is wrong-headed because it holds truth hostage to the existence of a maximally coherent framework. Assuming at the very least that the framework in question must belong to some human being(s) it seems unlikely that any human's framework is maximally coherent. This would mean no judgments are true.

A more plausible suggestion would be to say that

(CT): The moral judgment P is true if and only if P would be a member of a maximally coherent moral framework.

Read as a proposal for what makes moral judgments true (what manifests their truth) (CT) is attractive in part because it offers to secure some degree of Objectivity for moral truth. Not all moral judgments would be members of a maximally coherent system of such judgments. Consequently, (CT) would seem to avoid the moral triviality thesis according to which any old moral judgment can be true.

One might think that (CT) threatens to founder on the so-called conditional fallacy.¹⁶ This is a problem that can plague attempts to

¹⁵ This would make the position close to that advocated—on a global scale—by B. Blanshard. See his “Coherence as the Nature of Truth” from *The Nature of Thought*, vol. 2 (New York: Harper Collins, 1939), 260–79.

¹⁶ As an objection to epistemic accounts of truth, this point was first made by Plantinga (1982), who used it to slightly different ends; Wright (2001, 767, 2003, 120–2) generalizes it; and it is his account we draw on here.

define a categorical statement in terms of a subjunctive conditional. How the problem applies here can be brought out if we take the proposition in question in (CT) to be

(not-C): Maximally coherent systems don't exist.

Substituting (not-C) for "P" in (CT), we arrive at the conclusion (to put it intuitively) that maximally coherent systems don't exist just when some maximally coherent system would say they don't. And that makes little sense.

Our coherence theory is local, so the advocate of (CT) can object, reasonably, that (not-C) is not a moral judgment, and hence not open to inclusion into (CT). Fair enough. But this response only serves to highlight a less formal, but one might think, more formidable problem. Let us grant that the view does not require us to admit that maximally coherent frameworks actually exist. Nonetheless, it does require us to understand *what they would be like if they were to exist*. And one might question the account on this basis alone. That is, one might wonder whether we do have a sufficient grasp of the very concept of a maximally coherent framework. This misgiving emerges when we compare our use of maximal coherence in understanding warrant with our use of it in (CT). In the case of warrant, one might reply to our misgiving by noting in that case, "maximal coherence" merely names a vanishing point on the horizon towards which a system moves as it becomes *more coherent*. And we needn't have a clear grasp on what a maximally coherent system is like in order to understand what it would mean for a framework to become more coherent. To become more coherent is simply to have more of the coherence-making features to a greater degree. But in (CT) we no longer are simply talking about *approaching* an ideal limit. We are saying that what is true is so because it would be believed at that limit. So the question of what that limit is like is clearly legitimate.

These considerations suggest that we look elsewhere for an account of moral truth than (CT). But they also point us in a promising direction. Why not construct our notion of moral truth

from the materials that we *do* understand—from the idea of a framework becoming more and more coherent—or improving in coherence?

4. From Supercoherence to Concordance

A judgment is *superwarranted* when it has the property of surviving arbitrarily close examination without defeat. On a coherence theory of warrant like the one we've canvassed above, such examinations consist in seeing whether the moral judgment in question, based on evidence available to the ordinary reflective inquirer, is a member of a coherent moral framework, and would continue to be so under all increases of information, moral and non-moral.

This can be sharpened. In effect, we have already said what it is for a moral framework to increase or improve in coherence:

IMPROVEMENT: Framework F is more coherent than F* when F, on balance, has either more of the coherence-making features or some of those features to a greater degree.

This in turn can be used to define what it is for a given judgment to belong to a coherent framework. A judgment belongs to a coherent moral framework just when it is one of the judgments comprising that framework and it *coheres with* that framework:

WITH: P coheres with moral framework F if, and only if, the result of including P in F would, on balance, make F more coherent.

Given that completeness, consistency, and explanatory connectedness are coherence-making features, for example, adding a consistent and explanatorily connected judgment to the system will increase that system's coherence along those dimensions.

With these definitions in hand, we can now more precisely define what it is for a moral judgment to supercohere:

SUP: The moral judgment P supercoheres with F if and only if P coheres with F at some stage of inquiry and would continue

to do so without defeat, through all successive and additional improvements to F.

A moral framework improves in the sense defined when it grows more coherent. So a moral judgment, such as my judgment that torture is wrong, supercoheres with my moral framework when it would remain part of that framework without fail no matter how that system might improve—which is to say, no matter how much additional moral and non-moral judgments might be included in the system that increase that system’s coherence.

Even so roughly characterized, supercoherence seems a plausible candidate for the property that manifests truth for moral judgments. That is, while truth as such won’t be supercoherence, it might be that the property that manifests truth for moral judgments is supercoherence. Call a moral judgment supercoherent for short when it supercoheres with some moral framework. Thus given SUP, we can say that:

SC: For any moral judgment P, P is true if and only if P is supercoherent.

The advantage of (SC) over (CT) is that we don’t need to grasp at all of what a *maximally* coherent moral framework might be like.¹⁷ We need only grasp what makes a given framework coherent, what it is for a framework to improve in its coherence, and what it is for a belief to belong to that framework. Yet one of the things that makes supercoherence “super” is that moral judgments that have the property are, or would be, members of frameworks that are not

¹⁷ Of course, we can, if we wish, construct another conception of a maximally coherent system from our definition. Were all my moral judgments to supercohere, they would all cohere and continue to cohere, no matter how the system might be added to or improved in coherence. Thus maximal coherence, while idealized, can itself be built out of familiar, non-idealized materials: we can say that a maximally coherent system is the type of system that our judgments would compose were each individual judgment to have the categorical property of supercoherence, that is:

S is maximally coherent system if and only if all of its members would be supercoherent at some stage of inquiry and would remain so in every successive stage of inquiry.

merely static in their coherence but are improving in that regard.¹⁸ This drops right out of our above characterization of coherence and improvement. And this in turn shows that like (CT), (SC) clearly avoids moral triviality: not just any old moral judgment will be supercoherent. Supercoherence is a significant achievement for any moral judgment.

Yet (SC), while superior to (CT), is still inadequate as a theory of moral truth. For (SC), while avoiding moral triviality, is arguably still too permissive. Essentially, it says that a moral judgment is true when it would be part of a durable system of moral and *non-moral* judgments. Crucially, there is no requirement that the non-moral judgments themselves be *true*. And that is a problem. For it is likely that some absurd moral judgments will be durably coherent with a set of fixed but false non-moral beliefs about the world. Consider a misogynistic culture which deprives women of rights partly on the basis of a whole range of mistaken non-moral beliefs (e.g. some mistaken views intelligence or some other non-moral matter, together with mistaken views on what would count as evidence for views about those matters and so on). In other words, if certain false non-moral beliefs are forever held fast, and enough other adjustments are made to the system to compensate for holding them fast, even the craziest moral views might turn out to be supercoherent, and thus even the craziest *moral* views might be true.

The suggestion I take from this is that if it is to be true, it is not enough for the fabric of our moral thought to be woven tightly—to be durably coherent—it must also be nailed down, or grounded on a firmer floor. Otherwise, the continual improvement of a supercoherent framework might turn out to be improvement in the wrong direction.

Arguably, the structure of the functionalist theory itself allows us to provide further grounding for our moral thought. As we noted in Chapter 2, any such foundation would be unavailable were (SC)

¹⁸ That is, it will continually improve provided it continues to incorporate new judgments that increase the coherence of the system.

being proposed as a global theory of truth. I argued there that, where “F” names some moral framework, claims of the form

(A) F is coherent

or

(B) J coheres with F,

cannot themselves be true in virtue of cohering with F on pain of absurdity. In the present context, this point comes to this. We can insist that what makes a given moral framework coherent—what makes it have the coherence-making features to the degree that it has, in other words—is not a matter for *that* system itself to decide.¹⁹ But of course the functionalist can accept this result, since she is not advocating supercoherence as the nature of truth. She is claiming that it is simply one possible manifestation of truth. And this means that we can affirm what stands to reason in any event: Propositions like (A) and (B) are not moral judgments. (A) concerns the wholly different subject of whether a given system of moral judgments has a sufficient amount of the coherence-making features to a sufficient degree. And whether this is so is not a moral matter. Therefore, whatever makes moral judgments true will not be what makes propositions like (A) true.

Of course, it is an interesting—if independent to our present concern—question as to what *does* make propositions like (A) or (B) true. It would be a severe mistake to hold that claims about coherence form a single domain, as I earlier defined “domain”. That is because, as I argued above, “coherence” is best understood as a label for a—possibly open-ended—*collection* of features. And these features are themselves quite disparate in nature: ranging from the purely logical (“consistency”) to the aesthetic (“simplicity”) to purely epistemic (“explanatory”). Aesthetic and logical concepts arguably do compose distinct domains. They are notoriously *sui generis* in their nature—distinct and unique ways of classifying and

¹⁹ The requirement does not rule out that judgment (A) might be true in virtue of being a member of some *other* coherent framework of course.

evaluating the world. Epistemic concepts on the other hand, might themselves be quite multifarious in their nature. In any event, it is not my current project—nor, I think, could it be a project for a single book—to say what manifests truth for every kind of proposition. Here, I am only making the point that functionalism allows us to say what I suspect anyone attracted to a coherence theory of moral truth has essentially wished to say: namely that it is not a moral matter whether a given moral judgment is coherent or a member of a coherent system. And this fact allows us at least part of the needed ground: for whether a given morality is coherent will not be something determined by that morality itself.

More generally, the functionalist can insist—as again adherents of more traditional coherence theories could not—that what makes non-moral judgments true is not supercoherence. We can add a further constraint to a broadly coherence theory of moral truth that makes it far less permissive and therefore more plausible. That is, we can say that if membership in a durably coherent system is going to make a judgment true, that system must not only be *internally* supercoherent, it must be durably coherent with the *external* coherence-independent facts—with whatever kinds of judgments are true, in other words, by virtue of corresponding to an extra-human reality. Thus no moral framework, no matter how internally and durably coherent in the face of criticism, should convince us that gender affects intelligence, or that AIDS can be caught through sweat, or, for that matter, that some system is coherent when in fact it is not.

We can say that a framework that meets both of these constraints—internal supercoherence and durable continual coherence with the facts—is composed of *concordant judgments*. *My suggestion, in short, is that we see our moral inquiries as aiming at constructing frameworks of concordant judgments*. Such systems, were there ever to be any, would be durably improving coherent frameworks of judgments, some of which—the non-moral judgments—are true in virtue of corresponding to the facts, but others of which—the moral judgments—are true by supercohering to that very framework,

that is, by durably belonging to the framework itself. Thus, where P is a moral judgment,

(CM) P is true if, and only if, P is concordant.

And

P is concordant if and only if, (a) P supercoheres with a moral framework; and (b) that framework's morally relevant non-moral judgments are true.

Thus for a moral judgment to be true, it is not enough for it to be supercoherent, it must also be durably coherent with the relevant *non-moral truth*—that is our moral judgments must remain coherent in the face all of future increases of factual information about the relevant non-moral world.²⁰ Consequently, our hypothetical misogynistic judgments above would be ruled out, for it is simply not true that gender effects intelligence. Therefore, the judgment that women lack certain rights because they are less intelligent or some such would not be concordant with the moral and non-moral truths, and so false.

Assuming as we are the functionalist theory of truth as a background, the analysis given in (CM) and our definition of concordance is not circular. For the functionalist will claim that all *moral* judgments are true because they are concordant, non-moral judgments will be true in virtue of some other manifesting property. Moreover, truth *itself* is defined functionally.

I began this chapter for making a *prima facie* case for two claims: that moral judgments can be true, and yet their truth is not plausibly manifested by correspondence with a mind-independent

²⁰ What isn't relevant? Well, presumably any non-moral fact that humans couldn't know to obtain. After all, if no one could ever know whether p, its being the case that p, could hardly be part of the explanation for why some act is blameworthy or praiseworthy. This is why we find it odd to think that some action could be wrong because of some fact about the world that only God could ever know. If we can't ever know why an action is wrong, it is hard to see how we could be blamed for doing it. And that makes us (or should make us) question whether it is wrong in the first place. See the Appendix to this chapter for further discussion.

reality. Moral truth, I argued, does not seem to be a matter of correctly representing human actions as having special sorts of moral properties. Moreover, and as I will argue in the Appendix, the moral domain fits the three constraints a domain must meet if it is to have the truth of its judgments manifested by an epistemically-constrained property like superwarrant.

The resulting picture of moral truth is not unfamiliar: it sits squarely within the so-called constructivist tradition, which takes moral truth to be constructed out of the rational norms that govern moral judgment. Those norms, I've claimed, tell us that warrant for our moral judgments is a matter of those judgments having what I called "coherence-making" features. And thus the property that manifests truth for moral judgments—what I called concordance—is constructed out of those features. Yet concordance isn't identified with truth as such. Rather it is a property that manifests truth.

According to alethic functionalism, a property manifests truth by playing the truth-role. And a property plays the truth-role when it has the truistic features marked out by our core truisms. Concordance would seem to satisfy this constraint. First, given the principle that

P if and only if P is concordant

(CM) is consistent with the T-schema. Second, concordance is distinct not only from mere warrant, it is distinct even from superwarrant: a judgment can be superwarranted by the evidence at some stage of inquiry but not concordant. Thus the possibility for moral error is ample: just because I and everyone else I know justifiably believe some proposition doesn't mean that it is concordant. Likewise concordance is a stable notion: if a proposition is concordant, then it is concordant at any stage of inquiry. And finally concordance has the relevant normative import: it could hardly fail to be correct to believe what is concordant, and it seems plausible, given the Rawlsian account of moral theorizing, that we aim at concordant judgments when we engage in moral inquiry.

The concordant theory of what manifests moral truth is ecumenical. It is consistent with typical Humean and Kantian approaches to moral theory. Kantian metaethics is broadly Apollonian in approach: it supports a cognitivist, rationalist approach to moral judgment according to which such judgments can be true or false. This our account of moral truth clearly does as well. Humean approaches on the other hand, as Blackburn has emphasized, tend towards more Dionysusian virtues.²¹ They take ethics not to be just a matter for reason, but as emerging out of our feelings, our sentiments, our all too human passions. As we noted above and in the last chapter, such approaches typically reject the idea that moral judgment represents facts, and on that basis have denied that moral judgments have been true or false. But this is a mistake. With Blackburn, Gibbard, and Hume, we can take moral judgment to function primarily as an avenue for expressing our moral sentiments, but still hold that these judgments are capable of getting things right and getting them very wrong. Moreover, we can do this in a way that still allows us a strong contrast between the truth of moral judgments and the truth of judgments about the physical world. In short, the above theory opens up the possibility that we can be expressivist about moral judgment but cognitivist about moral truth.

Does taking concordance as manifesting truth require a revision of classical logic in domains where it does so? Put another way, is the advocate of concordance as an account of what manifests moral truth committed to what we earlier called domain-specific logical pluralism? That depends on whether we think it plausible that there are some propositions which are such that neither they nor their negation are concordant. Consider, to use an example employed earlier in another context: propositions like

(S) Sophie's choice is morally right.

Can we be sure that such a proposition either is or is not concordant? So long as we don't think there is any guarantee that every

²¹ Blackburn, *Ruling Passions*.

moral judgment is either concordant or not concordant, we would have warrant for thinking that bivalence, and more generally LEM, fail for the moral domain.

What would explain such a lack of confidence? Presumably it would be the thought that there are or could be some moral judgments such that adding either that judgment or its negation would fail to improve the system in the sense of “improvement” defined above. This might occur for trivial reasons if the proposition is simply completely unrelated, or is patent nonsense. Such propositions, according to our above definitions, would fail to be concordant because they would fail to supercohere. Adding them to our moral framework fails to increase, on balance, the coherence-making features of the framework. But it might also be that adding some propositions or their negations might actually make the framework *less* coherent. This would be the case if one thought that some judgments are so destructive morally speaking that they simply wreak explanatory havoc on the rest of our moral thought. (S) might be such a proposition. Sophie’s choice, in the novel by Styron, is to give one of her two small children away to the Nazi camp guard to be killed. She must give one of them to the guard or both are taken. (S) claims that her actual choice—to give up the younger of the two children—is the morally right one. But it is difficult to see how that could be. The more natural thought is that neither choice is right; hence we have warrant to think that any judgment to the contrary is neither concordant nor not concordant, and hence neither true nor false.

In sum, the concept of concordance all by itself does not *demand* a revision of classical logic, but it is consistent with it. And importantly, it does allow us to explain how we would be warranted in thinking that a judgment like (S) is neither true nor false.

Obviously, our account of concordance does not constitute a complete metaethic. But if this sketch of a theory is on the right track, it demonstrates the power of the functionalist theory of truth when combined with the metaphysical thesis of pluralism as I’ve understood that thesis in this book: that there is more than one

property that can manifest truth. For with functionalism as our background theory of truth, we can make sense of the idea that moral judgments are truth-apt, without committing us to cluttering the world with a new special sort of property. Moreover, it allows us to do this without committing the serious blunder of identifying truth with a property like concordance across the board.

6. Concordance, Relativism, and Skepticism

The concordance theory of moral truth is pitched at a certain set of intuitions about morality. Those intuitions steer us away from thinking that moral judgments are true in virtue of representing mind-independent moral properties. But they also steer us away from thinking that moral judgments are mere undisciplined expressions of emotions, not subject to rational appraisal. The concordance theory supplies an answer to how moral judgments can be true that attempts to satisfy both intuitions. It attempts to get the degree of Objectivity appropriate to morality right.

But does it? Some will worry that concordance is not objective enough. They'll ask: what rules out the possibility that there might be two concordant but inconsistent judgments of some action? That is, might not there be more than one view about some moral matter that would hang on equally well in the face of all future criticism, and accord equally well with the relevant empirical facts both now and forever into the future? And might not these be inconsistent to boot?

These questions are reasonable. People often seem to have different moral sensibilities; they judge character and action along distinct lines and standards. And even when people initially share the same moral sensibilities, those sensibilities can diverge in the face of differing experiences. Moreover, one might think that the very fact that the "is" of reality does not entail the "ought" of morality should tell us that the requirement that a concordant judgment cohere with the non-moral truth is not going to prevent more than one set of moral judgments from meeting that requirement.

And that alone might make us worry that the concordance theory is still too permissive, as I put it earlier.

At rock bottom, the worry here is that the advocate of the concordance theory of moral truth is a closet relativist. Admittedly, not an “anything goes” relativist; nor a “normative relativist” who stupidly counsels us to never criticize another’s moral judgments. But a relativist in the sense of admitting that there may be more than one true moral story of the world.

This worry cannot be ignored; but it can be de-fanged. It helps to first distinguish what we might call its practical face from its theoretical face. The practical face of the problem emerges when we consider the possibility of *encountering* an equally concordant but divergent moral framework.

So let’s consider the admittedly odd situation where we have *already* granted that someone else has an equally admirable and concordant set of moral opinions that are nonetheless inconsistent with our own. We presumably have evidence for thinking that

- (1) J is concordant with F
- (2) \sim J is concordant with F*

Suppose my framework is F. But I think that F* is equally admirable. That is, I also think that, with regard to J-relevant matters,

- (3) Neither F nor F* are capable of improvement.

It is not at all clear that this is a stable epistemic position.²² For recall how difficult it is for my judgments to achieve concordance. Coherence itself is difficult to achieve: it is not mere consistency as I’ve been at pains to show. Moreover, concordance requires durable coherence not only with the non-moral facts, but with other moral judgments in the face of all future improvement to those judgments—that is, in all stages of inquiry. And of course, *those future improvements must include whatever judgments are the result of an encounter with a divergent moral framework.* Accordingly it is

²² As Blackburn has argued; see his *Spreading the Word* (Oxford: Oxford University Press, 1984), 200–2.

difficult to see how I could have warrant for (3)—for it to cohere with my framework—and maintain warrant for both (1) and (2). The situation where *I* come to grant that $\sim J$ is concordant with F^* and F^* is incapable of improvement is a new piece of information for me. It is hard to see how this new fact wouldn't force me to alter my judgments in some way. That is, it seems I must either continue to hold J , and thereby lessen my admiration for F^* ; change my mind in favor of $\sim J$; or reject them both and hold that they are only partial reflections of a fuller story. The third option seems most likely. As Blackburn remarks, “in so far as acquaintance with another value system makes me respect it, then it properly makes me rethink both systems”.²³ In other words, when I come to recognize that my moral judgments and another's inconsistent judgments are equally warranted, I ipso facto have warrant for thinking that something has gone wrong with the debate. I may feel that the terms of the debate are too simplistic and need to be rejected; or I may, as in the Sophie's choice situation, think that neither J nor $\sim J$ improve either framework's coherence. We are damned if we do and damned if we do not.

The same point works in the other direction. If the judgment J is concordant then I am warranted in making that judgment. And if I am warranted in judging J then I am presumably warranted in judging that it is not warranted to judge not- J . So if, from my present stage of inquiry, J is concordant, then I am warranted in judging that it is not warranted to judge not- J . Hence where J is concordant, I will never be warranted in judging that both J and its negation are equally warranted. And if I can't grant they are equally warranted, then it is hard to see how I can grant that neither system is capable of improvement. And so it is hard to see how I could judge that both judgments are equally concordant.

The overall lesson here is this. If moral truth is concordance, it is difficult to see how I could be warranted in granting (1), (2) and

²³ As Blackburn has argued, *Spreading the Word*, 201.

(3). So I will never be in a position where I will be warranted in thinking that someone else's divergent moral framework is equally concordant as my own. If another's view seems worthy of respect, then I should either revise my own opinions, or, should both views seem equally worthy, regard neither as concordant, and hope that the truth will emerge in due time.

Still, the objector may feel that the real question has been dodged. Might not it simply *be* the case that there are two divergent frameworks each of whose judgments is equally concordant? Such a situation seems logically possible, even if we could never recognize it should it obtain. You and I may have inconsistent but concordant judgments even if we would never be warranted in judging that we do.

This is logically possible. But I don't see this mere possibility as a vice, but as a virtue. The fact that there may be more than one true story of the world when it comes to how to live a good life is hardly news (except maybe to some philosophers). The important thing is that given the constraints on concordance, not every story is equally true; many will never make the cut.²⁴ Consequently, in my view concordance supplies us with an objective enough notion of moral truth.

Others will worry that concordance, far from being not objective enough, is too objective. It only takes a moment to see that many of our moral opinions are probably not going to have the exalted status of concordance. Indeed, most of what we say about morality is probably not even coherent, let alone supercoherent. We have prejudices we are unaware of, and cause harm that we should

²⁴ Won't this entail that concordance is a relative notion? The definition itself entails that it is relative in the sense of being a relation between judgments and frameworks. But is it possible that the judgment that *p* might be concordant relative to one framework but not to another, and hence that moral truth is itself relative? It certainly raises this as a possibility. But as argued above, the ensuing relativism, has no epistemic or normative consequences. So it seems far from pernicious. Indeed, as noted in the text, it secures for us the intuition that in principle at least, there is more than one true moral story of the world. For further discussion, see Capps, Lynch, and Massey, "A Coherent Moral Relativism", *Synthese*, in press.

not. Won't the difficulty in achieving concordance mean that skepticism about our own moral opinions is merited?

I find it hard to sympathize with the complaint that the above theory makes moral knowledge difficult to achieve. It *is* difficult to achieve. Knowledge requires truth. And truth, like most things worth having, is hard to come by. It is all the more precious for being so.

Appendix: Representations and Epistemic Constraint

This chapter has made the case for the idea that a property constructed out of our epistemic norms for morality could serve to manifest truth for our moral judgments. As such it makes room for a particular possibility. But what reasons might we have for thinking this possibility is actual—for thinking, in short, that moral truth really isn't manifested by something like correspondence? Above, I pointed to one such reason. Namely, if we understand correspondence in terms of representation, and representation in the way I have in this book, it is difficult to see how moral judgments could correspond. In this brief appendix I want to add two further considerations in favor of that conclusion.

The first is that there is good reason to think that moral judgments aren't in the business of representing in any case. So if correspondence is a matter of accurate representation, moral judgments can't correspond to reality.

We'll have to come at the point the long way around. Expressivists have long argued that moral judgments aren't *beliefs*. And one might think this point, if sound, might be sufficient to establish that moral judgment is not representational. It is not sufficient, but it is worth pausing to see why. One way of formulating the classic expressivist argument, due to Shafer-Landau, is as follows:

1. Necessarily, if S judges that x is right, then S is motivated to some degree to act in accordance with that judgment.
2. Taken by themselves, beliefs do not motivate or generate any motivating states.

Therefore, moral judgments are not beliefs.²⁵

²⁵ See R. Shafer-Landau, *Moral Realism: A Defence* (Oxford: Oxford University Press, 2003), 121; versions differ, but same basic line of inference can be found in C. L. Stevenson, *Facts and Values* (New Haven, CT: Yale University Press), 14–31; see also R. M. Hare, *The Language of Morals* (Oxford: Oxford University Press, 1952), 79–93, and Blackburn, *Spreading the Word*, 188–9.

Both premises are alleged to be conceptual truths. The first premise is a form of what is called motivational judgment internalism.²⁶ The basic thought here is that morality is essentially practical. To judge that something is right or wrong is to be moved to act. Hence the internalist takes there to be a conceptual connection between moral considerations and motivations for action. And indeed there is much plausibility to this thought: when someone is not willing to act on their moral judgments, we typically doubt the sincerity of those judgments.

The second premise comes from what is sometimes called Humeanism about motivation. According to this view, the fact that I believe that someone is in danger and that it is right to help people in danger, does not automatically mean that I will help. I also must want to help them. It is desires that motivate, not beliefs all by themselves.²⁷ One reason for thinking so, highlighted in recent years by Michael Smith, is that beliefs have the wrong *direction of fit* to be motivational states. To be motivated is to have a goal; and to have a goal is to be in a state like a desire, that the “world must fit”. Beliefs, the thought runs, have their direction of fit the other way around. Beliefs aim to fit the world; desires aim for the world to fit them.

Given these two premises, some philosophers have concluded that moral judgments should not be understood as beliefs. This has led others to further conclude that moral judgments can't be true or false.²⁸ But the motivation argument does not in fact warrant either conclusion. To even reach the conclusion that

²⁶ This term comes from S. Darwall, *Impartial Reason* (Ithaca, NY: Cornell University Press, 1983), 51; the view can come in a variety of strengths, and should of course be distinguished from reasons internalism; for discussion see Shafer-Landau, *Moral Realism*, 143–5 and M. Smith, *The Moral Problem* (Oxford: Basil Blackwell, 1994), ch. 3. For a recent attack against all forms of internalism, see Bloomfield, *Moral Reality*.

²⁷ See D. Stampe, “The Authority of Desire”, *Philosophical Review*, 96 (1987), 335–81; Davidson, “Intending” in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980); Smith, *The Moral Problem*.

²⁸ Again, Stevenson, *Facts and Values*, 14–31; Hare, *The Language of Morals*, 79–93; and Blackburn, *Ruling Passions*, 61–5 is also sympathetic.

moral judgments are not true or false, one would have to further assume that

Only beliefs (or believed propositions) can be true or false.

And this might be contested. But even if we grant that point, it still doesn't follow that moral judgments aren't beliefs. As we noted in the last chapter, it is plausible that to qualify as a belief, a mental state needn't aim at representing the facts; it merely must aim to be true. These are equivalent only if one assumes that to be true is to (correctly) represent the facts. And that is an assumption a functionalist will reject.

That in turn means that one can accept something like the motivation argument, accept too that only beliefs can be true or false but reject the representational characterization of belief. Indeed, one might think that the reason given above for accepting premise 2—namely that beliefs have the wrong direction of fit in order to be motivating states, is mischaracterized. When we think of a belief fitting the world we are thinking of it *as a representational state*. But if, as the thin account of belief allows, there can be beliefs that are not representational, it needn't follow that moral judgments have the wrong direction of fit to be beliefs. It only follows that moral judgments can't be representational beliefs.

So what the motivation argument shows, at best, is not that moral judgments can't be true or false, or that they aren't beliefs, but that they aren't mental representations.²⁹ And that in turn demonstrates that those who take moral truth to be manifested by correspondence face a heavy theoretical burden. At the very least, they must not only defend the claim that moral judgments express or are beliefs, but, in addition, they must hold that moral judgments are by their nature representational. If moral judgments are not representational, it is difficult to see how they can correspond to moral facts. Consequently, it is a clear point in favor of a theory of

²⁹ Shafer-Landau, *Moral Realism*, 119–41, offers a particularly exhaustive discussion of the ins and outs of the debate.

moral truth if it can explain how moral judgments can be true even if moral judgments should turn out not to be representational. That is, it is a point in favor of a view of moral truth if it is consistent with the view that, as Mark Timmons puts it:

A moral judgment... is a certain contentful psychological state that is implicated in a web of defeasible psychological tendencies aimed primarily at choice and guidance of action, not representation.³⁰

The theory that moral truth is manifested by concordance is consistent with just this position on moral judgment.

A second reason to think that moral truth is not manifested by correspondence is that *moral truth is epistemically constrained*. Many writers, holding quite different normative ethical theories, think moral wrongness is conceptually tied to responsibility and blameworthiness. “We do not call anything wrong,” as Mill writes, “unless we mean to imply that a person ought to be punished in some way or other for doing it; if not by law, by the opinion of his fellow creatures; if not by opinion, then by the reproaches of his own conscience”.³¹ What is wrong is what we can be justifiably held responsible for doing—what is worthy of blame, in short. If so, then the following argument seems cogent:

If an action is wrong, then it is blameworthy.

If an action is blameworthy then it is feasible to know that it is wrong.

Therefore, if an action is wrong, it is feasible to know that it is wrong.

The argument is valid. The first premise is Mill’s conceptual truth. The second premise, too, seems to be a conceptual truth about

³⁰ Timmons, *Morality without Foundations*, 143. Timmons own view is that moral judgments can be treated as beliefs, even if they are not representational.

³¹ Mill, *Utilitarianism*, ch. v. See also S. Darwall, “Psychology and the Second-Person Standpoint”; R. Adams, *Finite and Infinite Goods* (1999), 238; K. Baier, “Moral Obligation,” (1966); R. Brandt, *A Theory of the Good and the Right* (1979), 163–76; A. Gibbard, *Wise Choices, Apt Feelings* (1990), 42; J. Skorupski, *Ethical Explorations* (1999), 142.

blameworthiness. It says that it is a necessary condition of being a blameworthy action that it is knowably wrong. The conclusion is an instance of what we earlier called an epistemic constraint:

If P then it is feasible to know P.³²

The intended modality is meant to range across truth-values in the actual world.³³ Applied in the present case, it is feasible to know that an action is wrong just when at least one human could at one time judge it wrong when it is.

The force of the argument lies in its simplicity. My action is blameworthy when it is possible for me to be held accountable for that action. But I can't be held accountable, surely, for an action that *no one, including myself*, would ever know is an action for which I should be censured.³⁴ Given the first premise, to think otherwise would imply not only that there are unnoticed and unknown moral wrongs but that there are unknowable moral wrongs. It might, for example, be wrong that I used more than six words in this sentence. And that seems absurd. Hence the conclusion seems warranted, and so, presumably would be a parallel conclusion about praiseworthiness and rightness. If so, then, thanks to the schema that it is true that p if and only if p, we know that there are no unknowable truths about what is morally right or wrong.

Note that the epistemic constraint on moral truths makes them noticeably different from truths about the natural world, where we are quite willing to accept unrecognizable truth. Facts about the inner life of bats, or what is outside our light cone seem like perfectly good facts, despite the seeming impossibility that any human will ever be able to recognize that such states of affairs obtain when they do obtain. Similarly, *mutatis mutandis* for mathematical truth.

³² For discussion of such principles, formulated in this way, see C. Wright, "On Being in a Quandary", *Mind*, 110 (2001), 461 ff.

³³ In contrast say, to propositions that are true in some non-actual possible world, and knowable as true in that world.

³⁴ Of course, not all will agree. Some act-utilitarians, for example, may not. It is worth emphasizing that many utilitarians will accept the second premise however: see notably the quote from Mill above.

An epistemic constraint on moral truth doesn't, all by itself, force one to give up on the idea that moral judgments are true just when they correspond to human independent moral states of affairs.³⁵ One might, for example, still insist that there are such states of affairs, but claim that it is simply a surd fact about the world and our cognitive capacities that we happen to be able to know when they obtain. Lucky devils that we are, we are just built to be able to discover such facts. Alternatively, one might say that fortune smiles not so much on us but on the moral properties themselves—moral properties, like colors and smells, are just the sort of properties that are recognizable by us. This is the position favored by some secondary property theorists and other “ontological constructivists”.³⁶ I leave it to others to determine how plausible these reactions are. For my part, I am inclined to consider them too convenient by far; I see no reason to think the world is so cooperative. I find it more plausible that the above argument does not tell us how about moral *reality* so much as it tells us about our judgments about morality. In particular it tells us something about how those judgments manifest truth.

³⁵ As Shafer-Landau rightly notes in, *Moral Realism*, 17.

³⁶ See especially J. McDowell, “Values and Secondary Qualities”, in T. Honderich (ed.), *Morality and Objectivity: A Tribute to John Mackie* (London: Routledge & Kegan Paul, 1985); and D. Wiggins, “Truth, Invention and the Meaning of Life”, in his *Needs, Values and Truth* (Oxford: Basil Blackwell, 1987).

Conclusion

The aim of this book has been to articulate and defend what I've called the functionalist theory of truth. That theory is motivated in part by the suspicion that if we are to ever come to grips with both the cognitive unity and semantic diversity of our thought, we need a new way of thinking about truth. We need a new theory, I've argued, because traditional theories lack the proper scope to account for diversity, while simple pluralist theories give up unity. And deflationism, while initially promising, is ultimately unsatisfying; among other problems, it removes truth from our explanatory resources.

The functionalist theory has two major components. The first is a functionalist elucidation of the concept and property of truth. It proceeds by defining truth by way of the role that it plays within our cognitive life. And it defines that role by appeal to the core truisms about truth—folk beliefs that constitute truth's nominal essence. The functionalist takes these truisms to also specify truth's real essence. This amounts to saying that truth just is the property that has the truish features essentially. Thus, like traditional theories, the functionalist believes truth is one.

The second major component of the functionalist theory is a detailed metaphysical account of how this single property can be multiply manifested. This is meant to capture the key intuition of pluralism—that truth is many. This intuition is preserved by functionalism, I've argued, because functionalism allows that truth is an immanent property. Where such a property is manifested by some other property, its essential features are included among the latter's features. This is what happens when truth is manifested by a correspondence property or the property of concordance. Such properties may manifest truth in some domains but not others.

When they do so, they play the truth-role by way of including the tru-ish features among their own features.

The overall picture of truth given here is abstract. But of course, that is part of the point. This book has been about truth itself—and truth itself, the functionalist theory claims, is a rather abstract property. But it is also a property that comes in more than one form. If that is right, then philosophical progress lies in investigating those forms, in discovering how truth manifests itself across the spectrum of our thought.

Select Bibliography

- Alston, William P., "Ontological Commitments", *Philosophical Studies*, 9 (1957), 8–17.
- *A Realist Conception of Truth* (Ithaca, NY: Cornell University Press, 1996).
- Aristotle, *Metaphysics*, trans. C. Kirwan (Oxford: Oxford University Press, 1993).
- Barnard, Robert, and Terence Horgan, "Truth as Mediated Correspondence", *The Monist*, 89 (2006), 28–49.
- Bar-On., D., C. Horsick, and W. Lycan, "Deflationism, Meaning and Truth-Conditions", in Jc Beall and B. Armour-Garb (eds.), *Deflationary Truth* (Chicago, IL: Open Court, 2005).
- Beall, Jc, "On Mixed Inferences and Pluralism about Truth Predicates", *Philosophical Quarterly*, 50 (2000), 380–2.
- "Transparent Disquotationalism", in Jc Beall and B. Armour-Garb (eds.), *Deflationism and Paradox* (Oxford: Oxford University Press, 2005).
- and B. Armour-Garb, *Deflationary Truth* (New York: Open Court, 2004).
- and Greg Restall, *Logical Pluralism* (Oxford: Oxford University Press, 2005).
- B. Armour-Garb, and G. Priest, *New Essays on Non-Contradiction* (Oxford: Oxford University Press, 2004).
- Blackburn, Simon, *Spreading the Word* (Oxford: Oxford University Press, 1984).
- *Ruling Passions* (Oxford: Oxford University Press, 1998).
- Blanshard, Brand, *The Nature of Thought*, vol. 2 (New York: Harper-Collins, 1939).
- Bloomfield, Paul, *Moral Reality* (Oxford: Oxford University Press, 2001).
- Boghossian, Paul, "The Normativity of Content", *Philosophical Studies*, 13 (2003), 31–45.
- *Fear of Knowledge* (Oxford: Oxford University Press, 2007).
- Brandom, Robert, "Explanatory vs. Expressive Deflationism about Truth", in R. Schantz (ed.), *What is Truth?* (Berlin: Walter de Gruyter, 2001).

- *Making it Explicit* (Cambridge, MA.: Harvard University Press, 1994).
- Brink, David, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989).
- Capps, D., M. P. Lynch, and D. Massey, “A Coherent Moral Relativism” *Synthese* (in press).
- Cotnoir, Aaron, “Generic Truth and Mixed Conjunctions: Some Alternatives”, *Analysis*, 69 (forthcoming).
- Damnjanovic, Nicolas, “Deflationism and the Success Argument”, *Philosophical Quarterly*, 58 (2005), 53–67.
- David, Marian, *Correspondence and Disquotation* (Oxford: Oxford University Press, 1999).
- “Minimalism and the Facts about Truth”, in R. Schantz (ed.), *What is Truth?* (Berlin: Walter de Gruyter, 2001).
- Davidson, Donald, *Inquiries into Truth and Interpretation* (Oxford: Clarendon Press, 1984).
- “The Structure and Content of Truth”, *Journal of Philosophy*, 87 (1990), 279–328.
- Devitt, Michael, *Realism and Truth* (Princeton, NJ: Princeton University Press, 1997).
- Dodd, Julian, “Recent Work on Truth”, *Philosophical Books*, 43:4 (2002), 279–91.
- Dretske, Fred, *Knowledge and the Flow of Information* (Cambridge, MA.: MIT Press, 1981).
- Dummett, Michael, *Frege: Philosophy of Language* (London: Duckworth, 1973).
- *Truth and Other Enigmas* (Cambridge, MA.: Harvard University Press, 1978).
- *The Seas of Language* (Oxford: Oxford University Press, 2003).
- Edwards, Douglas, “How to Solve the Problem of Mixed Conjunctions”, *Analysis*, 68 (2008), 143–9.
- Engel, Pascal, *Truth* (London: Acumen Press, 2002).
- Field, Hartry, “Tarski’s Theory of Truth”, *Journal of Philosophy*, 69 (1972), 347–75.
- “The Deflationary Conception of Truth”, in G. MacDonald and C. Wright (eds.), *Fact, Science and Morality* (Oxford: Blackwell, 1986), 55–117.

- “Critical Notice: Paul Horwich’s *Truth*”, *Philosophy of Science*, 59 (1992), 321–330.
- *Truth and the Absence of Fact* (Oxford: Oxford University Press, 2001).
- Fodor, Jerry, *Psychosemantics* (Cambridge, MA.: MIT Press, 1987).
- Fumerton, Richard, *Realism and the Correspondence Theory of Truth* (Chicago, IL: Rowman & Littlefield, 2002).
- Funkhouser, Eric, “The Determinable–Determinate Relation”, *Noûs*, 40 (2006), 548–69.
- Greenough, Patrick, and Michael P. Lynch, *Truth and Realism* (Oxford: Oxford University Press, 2006).
- Gupta, Anil, “A Critique of Deflationism”, *Philosophical Topics*, 21 (1993), 57–81.
- Harman, Gilbert, *The Nature of Morality* (Oxford: Oxford University Press, 1977).
- Hill, Christopher, *Thought and World* (Cambridge: Cambridge University Press, 2001).
- Horgan, Terence, “Contextual Semantics and Metaphysical Realism”, in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA.: MIT Press, 2001), 67–96.
- and Mark Timmons, “New Wave Moral Realism Meets Moral Twin Earth”, *Journal of Philosophical Research*, 16 (1991), 447–65.
- and Matjaz Potrč, *Austere Realism* (Cambridge, MA.: MIT Press, 2008).
- Horwich, Paul, *Truth* 2nd edn. (Oxford: Oxford University Press, 1998).
- *Meaning* (Oxford: Oxford University Press, 1999).
- “Norms of Truth and Meaning”, in Richard Schantz (ed.), *What is Truth?* (Berlin: Walter DeGruyter, 2001).
- Hyde, Dominic, “Pleading Classicism”, *Mind*, 108 (1999), 733–55.
- Jackson, Frank, *From Metaphysics to Ethics: A Defense of Conceptual Analysis* (Oxford: Oxford University Press, 1998).
- “Representation, Truth and Realism”, *The Monist*, 89 (2006), 50–62.
- James, W., *Pragmatism and the Meaning of Truth* (Cambridge, MA.: Harvard University Press, 1942).
- Kim, Jaegwon, *Mind in a Physical World* (Cambridge, MA.: MIT Press, 1998).
- Kirkham, Richard, *Theories of Truth* (Cambridge, MA.: MIT Press, 1992).

- Künne, W., *Conceptions of Truth* (Oxford: Oxford University Press, 2003).
- Lynch, Michael P., *Truth in Context* (Cambridge, MA.: MIT Press, 1998).
- “A Functionalist Theory of Truth”, in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA.: MIT Press, 2001).
- “The Truth in Contextual Semantics”, *Grazer Philosophische Studien*, 63 (2002), 173–95.
- “Minimalism and the Value of Truth”, *Philosophical Quarterly*, 54 (2004), 497–517.
- “Replies to Critics”, *Philosophical Books*, 46 (2005), 331–42.
- “Truth and Multiple Realizability”, *Australasian Journal of Philosophy*, 82 (2004), 384–408.
- *True to Life* (Cambridge, MA.: MIT Press, 2005).
- “The Values of Truth and the Truth of Values”, in D. Pritchard, A. Millar, and A. Haddock (eds.), *Epistemic Value* (Oxford: Oxford University Press, forthcoming).
- McDowell, John, “Values as Secondary Qualities”, in T. Honderich (ed.), *Morality and Objectivity: A Tribute to John Mackie* (London: Routledge & Kegan Paul, 1985).
- *Mind and World* (Cambridge, MA.: Harvard University Press, 1996).
- McGinn, Colin, *Logical Properties* (Oxford: Oxford University Press, 2001).
- McGrath, Matthew, “Lynch on the Value of Truth”, *Philosophical Books*, 46 (2005), 302–10.
- Mackie, J. L., *Ethics: Inventing Right and Wrong* (New York: Penguin, 1991).
- Millikan, Ruth, *Language, Thought and other Biological Categories* (Cambridge, MA.: MIT Press, 1984).
- Misak, Cheryl, *Truth and the End of Inquiry* (Oxford: Oxford University Press, 2004).
- Newman, Andrew, *The Correspondence Theory of Truth* (Cambridge: Cambridge University Press, 2002).
- Papineau, D., and G. MacDonald, *Teleosemantics* (Oxford: Oxford University Press, 2006).
- *Reality and Representation* (Oxford: Blackwell, 1991).
- Patterson, Douglas, “Deflationism and the Truth-Conditional Theory of Meaning”, *Philosophical Studies*, 124 (2005), 271–94.

- Pedersen, Nikolaj, "What Can the Problem of Mixed Inferences Teach us about Alethic Pluralism?", *The Monist*, 89 (2006), 103–17.
- Peirce, Charles, "How to Make our Ideas Clear", *Popular Science Monthly*, 12 (1878), 286–302.
- Pereboom, Derek, "Robust Nonreductive Physicalism", *Journal of Philosophy*, 99 (2002), 499–531.
- Pettit, Philip, "Realism and Truth: A Comment on Crispin Wright's *Truth and Objectivity*", *Philosophy and Phenomenological Research*, 56 (1996), 883–9.
- Plantinga, Alvin, "How to be an Anti-Realist", *Proceedings and Addresses of the American Philosophical Association*, 56:1 (1982), 47–70.
- Polger, Thomas, *Natural Minds* (Cambridge, MA.: MIT Press, 2004).
- Priest, Graham, *Doubt Truth to be a Liar* (Oxford: Oxford University Press, 2006).
- Putnam, Hilary, *Meaning and the Moral Sciences* (London: Routledge & Kegan Paul, 1978).
- *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981).
- *Realism with a Human Face* (Cambridge, MA.: Harvard University Press, 1990).
- Quine, W. V. O., *The Pursuit of Truth* (Cambridge, MA.: Harvard University Press, 1990).
- Rawls, John, "Kantian Constructivism in Moral Theory", *Journal of Philosophy*, 77 (1980), 515–72.
- Rorty, Richard, "Is Truth a Goal of Inquiry", in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA.: MIT Press, 2001), 259–94.
- Russell, Bertrand, *The Principles of Mathematics* (New York: Norton, 1903).
- "The Philosophy of Logical Atomism", in his *Logic and Knowledge: Essays 1901–1950* (London: George Allen & Unwin, 1956).
- "The Monistic Theory of Truth", in his *Philosophical Essays* (London: George Allen & Unwin, 1966), 131–146.
- "On the Nature of Truth and Falsehood", in his *Philosophical Essays* (London: George Allen & Unwin, 1966), 147–59.
- Sainsbury, Mark, "Crispin Wright: Truth and Objectivity", *Philosophy and Phenomenological Research*, 56 (1996), 899–904.
- Schiffer, Stephen, *The Things We Mean* (Oxford: Oxford University Press, 2003).

- Shafer-Landau, Russ, *Moral Realism: A Defence* (Oxford: Oxford University Press, 2003).
- Shah, Nishi, "How Truth Governs Belief", *Philosophical Review*, 112 (2003), 447–83.
- Shapiro, Lawrence, "Multiple Realizations", *Journal of Philosophy*, 97 (2000), 635–54.
- Shapiro, Stewart, "Truth and Proof: Through Thick and Thin", *Journal of Philosophy*, 95 (1998), 493–521.
- Sher, G., "In Search of a Substantive Theory of Truth", *Journal of Philosophy*, 101 (2004), 5–36.
- "Functional Pluralism", *Philosophical Books*, 46 (2005), 311–30.
- Shoemaker, Sydney, "Realization and Mental Causation", in C. Gillett and B. Loewer (eds.), *Physicalism and its Discontents* (Cambridge: Cambridge University Press, 2001).
- Smith, Michael, *The Moral Problem* (Oxford: Basil Blackwell, 1994).
- Soames, Scott, *Understanding Truth* (Oxford: Oxford University Press, 1999).
- Sosa, Ernest, "Epistemology and Primitive Truth", in M. P. Lynch (ed.), *The Nature of Truth* (Cambridge, MA.: MIT Press), 641–62.
- Strawson, P. F., "Truth", *Proceedings of the Aristotelian Society*, supp. 24 (1950), 129–56.
- Tappolet, Christine, "Mixed Inferences: A Problem for Pluralism about Truth Predicates", *Analysis*, 57 (1997), 209–11.
- "Truth, Pluralism and Many-Valued Logic: A Reply to Beall", *Philosophical Quarterly*, 50 (2000), 382–5.
- Tarski, Alfred, "The Concept of Truth in Formalized Languages", in A. Tarski, trans. J. H. Woodger, *Logic, Semantics, Metamathematics* (Indianapolis, IN: Hackett, 1983).
- Timmons, Mark, *Morality without Foundations* (Oxford: Oxford University Press, 1998).
- Velleman, David, *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000).
- Vision, Gerald, *Veritas* (Cambridge, MA.: MIT Press, 2004).
- Walker, Ralph, *The Coherence Theory of Truth* (London: Routledge, 1989).
- Wedgwood, Ralph, "The Aim of Belief", *Philosophical Perspectives*, 16 (2002), 268–29.

- Williams, Michael, "Do we (Epistemologists) Need a Theory of Truth?", *Philosophical Topics*, 4 (1999), 223–42.
- "On Some Critics of Deflationism", in R. Schantz (ed.), *What is Truth?* (Berlin: Walter de Gruyter, 2001).
- Williamson, Timothy, "Critical Study of Truth and Objectivity", *International Journal of Philosophical Studies*, 30 (1994), 130–44.
- Wittgenstein, Ludwig, *Tractatus Logico-Philosophicus*, trans. C. K. Ogden (London: Routledge Kegan Paul, 1933).
- *On Certainty* (Oxford: Blackwell, 1969).
- *Philosophical Investigations*, 3rd edn., trans. G. E. M. Anscombe (New York: Prentice Hall, 1973).
- *Notebooks 1914–1916*. (eds.) G. E. M. Anscombe and G.H. von Wright, (trans.) G. E. M. Anscombe (Chicago: University of Chicago Press, 1984).
- Wright, C. D. "On the Functionalization of Pluralist Approaches to Truth", *Synthese* 145 (2005): 1–28.
- Wright, Crispin, *Truth and Objectivity* (Cambridge, MA.: Harvard University Press, 1992).
- "Truth in Ethics", *Ratio*, 3 (1995), 210–26.
- "Truth: A Traditional Debate Reviewed", *Canadian Journal of Philosophy*, 24 (1999), 31–74.
- "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism", *Mind*, 110 (2001), 45–98.
- *Realism, Meaning and Truth* (Oxford: Blackwell, 2001).
- *Saving the Differences* (Cambridge, MA.: Harvard University Press, 2003).
- Yablo, Stephen, "Mental Causation", *Philosophical Review*, 99 (1992), 499–531.
- "Singling out Properties", *Philosophical Perspectives*, 9 (1995), 477–502.

This page intentionally left blank

Index

- Adams, R. 188 n. 31
alethic functionalism, *see* functionalist theory of truth
alethic pluralism 91, 94–6, 100–3; *see also* pluralism; truth, ‘pluralism about’
Alston, W. 8 n. 2, 10 n. 4, 111 n. 15, 112 n. 17
antirepresentationalism 36–49, 145, 164
Aristotle 8
Armour-Garb, B. 16 n. 16, 109 n. 10
Armstrong, D. 75 n. 4
Audi, R. 166 n. 12
Ayer, A. J. 163 n. 4
- Baier, K. 188 n. 31
Bar-On, D. 144 n. 17
Barnard, R. 23 n. 4, 157 n. 26
Beall, J. 16 n. 16, 38 n. 22, 56 n., 92, 93, 94 n., 96 n. 13, 108 n. 8, 109 n. 10
belief 22, 53 n. 3, 112, 115, 131, 135, 149–53, 185–8; *see also* truisms about truth, ‘Norm of Belief’; warrant; content, ‘truth-conditional theories of’
Blackburn, S. 21 n. 1, 150–3, 156 n. 24, 161 n. 1, 163, 178, 181 n. 22, 182, 185 n. 25, 186 n. 28
Blanshard, B. 46 n. 31, 169 n. 15
Block, N. 143 n. 16
Bloomfield, P. 161 n. 1, 186 n. 26
Boghossian, P. 11 n. 7
Bonjour, L. 166 n. 11
Boolos, G. 63 n. 17
Boyd, R. 162 n. 3
Bradley, F. H. 46 n. 31
Brandom, R. 107 n. 6
Brandt, R. 188 n. 31
Brink, D. 161 n. 1, 165, 166
Bush, G. W. 15
- Capps, D. 183 n. 24
cognitive unity vs. semantic diversity 1–3, 4, 5, 50, 155–7, 159–64, 191–2
coherence 3, 109
 as theory of moral truth 168–71
 as theory of warrant 39, 165–8
 nature of 166–8, 171, 174–5
coherence theory of truth 5, 39, 46
 ‘many systems objection’ to 44, 47
 see also coherence; supercoherence
coherentism vs. foundationalism in ethics 165–6
concepts 23, 24, 80
concordance
 as manifesting truth in morality 177–8
 defined 175–6
concordance theory of moral truth 175–84
conditional fallacy 37, 169–70
constructivism in ethics 177
content 30, 41, 50, 136–42
 conceptual role theories of 136–42
 truth-conditional theories of 137–40, 143–4, 145, 148
 verification theories of 139, 148
 see also propositions; meaning; truisms about truth, ‘Content’
correspondence 1–2, 3, 18, 21, 23, 24, 34–6, 67 n. 21, 109, 110, 140–1, 162–3
correspondence theory of truth 5, 18–19, 21–2, 76, 87–8, 141, 156–7
 and cognitive science 22, 31
Cotnoir, A. 55 n. 7, 86 n. 1, 99 n. 16, 104 n. 19
- Damjanovic, N. 107 n. 4, 121 n. 23, 123
Daniels, N. 165 n. 9

- Darwall, S. 186 n. 26, 188 n. 31
 David, M. 21 n. 1, 90 n. 5, 107 n. 5,
 109 n. 12
 Davidson, D. 28, 138 n. 12, 186 n. 27
 definite descriptions 61; *see also* truth,
 'concept of'
 deflationism about truth 4, 5, 18–19,
 49, 105–114, 144, 156, 163–4,
 191
 disquotationalism 108
 minimalism 59; *see also* pluralism,
 'Wright's reductive pluralism'
 denotation 24
 causal theory of 25–7, 31, 138–9,
 141, 147
 teleological theory of 26–7, 31,
 138–9, 141
 determinate/ determinable
 distinction 75
 Devitt, M. 23 n. 4, 26 n. 9
 dialetheism 16
 Dodd, J. 55 n. 7, 86 n. 1
 domain specificity 79–82
 Dretske, F. 26 n. 9, 138 n. 13
 Dummett, M. 11 n. 9, 41, 58, 110
 n. 14, 132 n. 3, 139 n. 15, 159
- Edwards, D. 55 n. 7, 66 n. 18, 74 n. 2,
 86 n. 1
 Einstein, A. 105
 Engel, P. 10 n. 6
 epistemic constraint 42–4, 52–3, 142,
 176 n. 20, 185–90
 error theory 2, 4, 35, 49, 163
 essence of truth 84, 117
 nominal 18, 51, 71, 74; *see also*
 nominal essence
 real 18, 51, 71, 74, 191
 evidence, *see* warrant
 expressivism 2, 4, 35, 49, 150, 163,
 178, 185–6
- facts 22, 23
 fictionalism 35
 Field, H. 23 n. 4, 106 n. 3, 108 n. 7,
 108 n. 8, 109 n. 11, 121 n. 23
 Fitch's knowability paradox 42, 44
 Fodor, J. 26 n. 9, 32 n. 15, 138
 n. 13
- Frege, G. 133 n. 5, 144, 148
 Fumerton, R. 34 n. 16
 functionalist theory of truth 3, 5, 6,
 70–8, 82–4, 87–91, 94–104,
 114–20, 133 n. 9, 140, 145,
 153–5, 157–8, 191–2
 Funkhouser, E. 75 n. 4
- Gibbard, A. 178, 188 n. 31
 Greenough, P. 104 n. 19
 Gupta, A. 109 n. 10, 109 n. 13
- Hare, R. M. 185 n. 25, 186 n. 28
 Harman, G. 137 n. 11, 161 n. 1, 162,
 166 n. 11
 Hill, C. 106 n. 2
 holism 137, 148
 Horgan, T. 23 n. 4, 66 n. 18, 157
 n. 26, 162 n. 3
 Horsick, C. 144 n. 17
 Horwich, P. 24 n. 7, 106, 107 n. 6,
 108 n. 9, 109 n. 11, 109 n. 13, 111
 n. 15, 111 n. 16, 113, 116
 Hume, D. 178
 Hyde, D. 103 n. 17
- idealism 40–1
 immanent properties, *see*
 manifestation
 intentionality, *see* content
 van Inwagen, P. 130 n. 1
- Jackson, F. 17 n. 18, 31 n. 13, 35
 n. 17, 125 n. 28
 James, W. 10, 21, 154
- Kaplan, D. 119 n. 22
 Kim, J. 62, 65, 66 n. 19, 120, 123
 n. 24, 123 n. 25
 Kirkham, R. 90 n. 5
 Kripke, S. 54
 Künne, W. 8 n. 2, 54 n. 4
- Lewis, D. 17 n. 18, 62, 119 n. 22
 Locke, J. 7, 107
 logical modesty 99–104
 logical pluralism 91–104
 domain-specific 94–104, 178–9

- Lycan, W. 144 n. 17, 166 n. 11
 Lynch, M. 10 n. 6, 11 n. 7, 13 n. 11,
 14 n. 13, 55 n. 7, 66 n. 18, 73
 n. 1, 86 n. 1, 155 n. 23, 157 n. 26,
 183 n.
- MacDonald, G. 26 n. 10
 McDowell, J. 190 n. 36
 McGinn, C. 107 n. 4
 McGrath, M. 111 n. 16
 Mackie, J. L. 163 n. 5
 manifestation 3, 5–6, 60, 63, 66,
 69–70, 74–8, 81–2, 83–4, 90–1,
 117–18, 154–5, 191–2
 Massey, D. 183 n. 24
 meaning 142–5; *see also* content;
 propositions; semantic
 functionalism
 Mill, J. S. 188, 189 n. 34
 Millikan, R. 26 n. 10, 32 n. 15, 138
 n. 14
 Misak, C. 38 n. 21
 mixed compounds, problem of 56–7,
 62–3, 67, 85, 86–91, 97–9
 mixed inferences, problem of 55–6,
 63, 65, 86–7, 97, 99
 moral skepticism 183–4
- naturalism 23, 26
 Neander, K. 32 n. 15
 Newman, A. 34 n. 16
 Nietzsche, F. 51, 59
 nominal essence 7–8
- Papineau, D. 26 n. 10, 138 n. 14
 Patterson, D. 144 n. 17
 Peacocke, C. 32 n. 15
 Pedersen, N. 55 n. 7, 61 n. 15,
 63 n. 17, 67 n. 20, 86 n. 1, 99
 n. 16
 Peirce, C. 12, 36
 perception 25, 42
 Pereboom, D. 75 n. 5
 Pettit, P. 54 n. 4, 125 n. 28
 Pilate 7
 plain truth, *see* truth, 'self-manifest'
 Plantinga, A. 37 n. 20, 169 n. 16
 Plato 7, 69
- pluralism:
 simple alethic pluralism
 (SAP) 54–9, 60, 86
 Wright's reductive pluralism 59–67
see also truth, 'pluralism about'
- Potrč, M. 66 n. 18
 Priest, G. 16 n. 16, 168 n. 14
 propositions 9, 14–15, 76, 78–82,
 115, 129–34, 155–7
 domain-specificity of 81–2
 functionalist theory of 131–2
 truisms about 131–2
see also mixed compounds, problem
 of; mixed inferences, problem
 of
- Putnam, H. 36–7, 121 n. 23
- quasi-realism, *see* Blackburn
 Quine, W.V.O. 108 n. 7
- Ramsey, F. 108 n. 8
 Rawls, J. 164–5
 Read, S. 94 n. 11
 realism and antirealism, *see* repre-
 sentationalism; antirepresenta-
 tionalism; truisms about truth,
 'Objectivity'; superwarrant
 realization, *see* manifestation
 reference 23, 145–9; *see also*
 denotation; representation;
 representationalism
 reflective equilibrium 164–5
 relativism in ethics 181–3
 representation 2, 22–6, 31, 33, 151–3,
 161–2, 185–90; *see also* content,
 'truth-conditional theories of'
 representationalism 22–32, 42, 51–2,
 53, 77, 138, 145; *see also*
 representation; correspondence
 theory of truth;
- responsiveness 32–4, 42, 141, 142, 161
 Restall, G. 38 n. 22, 92, 93, 94 n. 11,
 96 n. 13
 Rorty, R. 40 n. 24
 Rossberg, M. 104 n. 19
 Rumsfeld, D. 10 n. 3
 Russell, B. 22, 44, 48, 88, 133 n. 4
 Ryder, D. 31 n. 14

- Sainsbury, M. 54 n. 4, 55 n. 7, 86 n. 1
 Sayre-McCord, G. 165 n. 9
 Scheffler, S. 165 n. 9
 Schiffer, S. 107 n. 4, 132 n. 3, 133
 scope problem 4, 34–6, 48–9, 109, 139
 Scott, F. 143 n. 16
 Sellars, W. 111, 137 n. 11, 160
 semantic functionalism 134–6, 139–40, 157–8
 Shafer-Landau, R. 166, n. 12, 185, 186 n. 26, 187 n. 29, 190 n. 35
 Shah, N. 11 n. 7
 Shapiro, Larry 118 n. 21
 Shapiro, Lionel 104 n. 19
 Shapiro, S. 109 n. 11
 Sher, G. 34 n. 16, 156 n. 25
 Shoemaker, S. 75 n. 4, 75 n. 5
 Skorupski, J. 188 n. 31
 Smith, M. 186
 Soames, S. 61 n. 14, 90 n. 5
 Stalnaker, R. 133 n. 6
 Stampe, D. W. 26 n. 9, 138 n. 13, 186 n. 27
 Stevenson, C. L. 163 n. 4, 185 n. 25, 186 n. 28
 Strawson, P. F. 23
 Sturgeon, N. 162 n. 3
 Styron, W. 98, 179
 success argument 121–6
 supercoherence 39–40, 41, 44–9, 148, 171–3
 superwarrant 38–44, 48–9, 52–3, 60, 78, 140–1, 142, 145–9

 T-schema 9, 28, 59, 106–8, 111, 115
 Tappolet, C. 55 n. 7, 56 n. 8, 57 n. 9, 86 n. 1
 Tarski, A. 24
 theory of truth, nature of 7, 17–19, 51, 82–4
 Timmons, M. 150, 156 n. 24, 162 n. 3, 163 n. 6, 188
 true generalizations, problem of 57, 63
 truisms about truth 5, 13, 15–17, 17–19
 Bivalence 16, 93, 96, 103, 179
 Content 14, 28
 core, *see* truisms about truth, ‘End of Inquiry’; ‘Norm of Belief’; ‘Objectivity’
 End of Inquiry 11–13, 14, 36, 38, 40, 57–8, 66, 70, 111, 153–5
 Non-contradiction 16
 Norm of Belief 10–11, 13, 14, 15, 17, 29–30, 40, 57–8, 70, 71, 111–13, 153–5
 Objectivity 8–10, 11, 13, 15, 17, 21, 22, 28, 29, 31, 36, 40–1, 45, 66, 70, 71, 110–11, 115, 157, 163
 Transparency 15
 Warrant Independence 10, 15, 17, 59
 truth:
 as explanatory property 4–5, 65, 113–14, 119, 120–7
 as many 68, 72, 78, 191; *see also* truth, ‘pluralism about’
 as metaphysically transparent property 106–8, 116–17
 as one 3–4, 5, 21, 67, 78, 191
 concept of 6, 60–2, 63, 72, 84, 106–8, 111, 115–19, 122
 in ethics 1, 150–1, 159–64, 171–80, 186–7
 internal realist definition of 37
 nature of 3, 5, 6
 pluralism about 6, 16, 50, 53, 58–9, 82, 85–6, 99, 101, 109, 131, 179–80; *see also* functionalist theory of truth; manifestation
 pragmatist theory of 36
 self-manifest 90–1, 97
 see also superwarrant; coherence theory of truth; correspondence theory of truth; deflationism about truth; functionalist theory of truth; essence of truth
 universality of reason 86, 91, 101
 validity 55–6, 92
 classical 92–104
 constructivist 92–104

- Velleman, D. 11
- Vision, G. 21 n. 1, 34 n. 16
- Walker, R. 45 n. 30, 46 n. 31
- warrant 9, 12, 37–9, 115
 moral 164–8
- weak grounding principle 90–1
- Wedgwood, R. 11 n. 7
- Wiggins, D. 190 n. 36
- Williams, M. 107 n. 6
- Williamson, T. 42 n. 28, 57 n. 9
- Wittgenstein, L. 22, 50, 85, 88–9,
 129, 134, 158 n. 27
- Wright, C. 10 n. 4, 16 n. 15, 37 n. 20,
 38, 41, 42 n. 27, 43, 52, 53 n. 2,
 57 n. 9, 73, 95, 115, 120, 139
 n. 15, 149, 161, 162, 166, 169
 n. 16, 189 n. 32; *see also* pluralism,
 ‘Wright’s reductive pluralism’
- Wright, C. D. 17 n. 17
- Yablo, S. 75 n. 5, 124