

Hard Truths

Elijah Millgram

 **WILEY-BLACKWELL**

A John Wiley & Sons, Ltd., Publication

Hard Truths

Hard Truths

Elijah Millgram

 **WILEY-BLACKWELL**

A John Wiley & Sons, Ltd., Publication

This edition first published 2009

© 2009 by Elijah Millgram

Blackwell Publishing was acquired by John Wiley & Sons in February 2007. Blackwell's publishing program has been merged with Wiley's global Scientific, Technical, and Medical business to form Wiley-Blackwell.

Registered Office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, United Kingdom

Editorial Offices

350 Main Street, Malden, MA 02148-5020, USA

9600 Garsington Road, Oxford, OX4 2DQ, UK

The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, for customer services, and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com/wiley-blackwell.

The right of Elijah Millgram to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Millgram, Elijah.

Hard truths / Elijah Millgram.

p. cm.

Includes bibliographical references and index.

ISBN 978-1-4051-8815-9 (hardcover : alk. paper)

1. Truth. 2. Reasoning. 3. Metaphysics. I. Title.

BD171.M49 2009

121-dc22

2008040258

A catalogue record for this book is available from the British Library.

Set in 10.5/13pt Minion by SPi Publisher Services, Pondicherry, India

Printed in Singapore by Utopia Press Pte Ltd

Contents

<i>Acknowledgments</i>	viii
1 Introduction	1
Part I Motivation	15
2 The Truth in Bivalence	17
3 Deflating Deflationism	33
Part II Arguments	49
4 How to Find Your Match	51
5 Unity of the Intellect	70
6 How Can We Think about Partial Truth?	102
Part III The Competition	123
7 Logics of Vagueness	125
Part IV Applications	147
8 The Quinean Turn	149
9 The Davidsonian Swerve	165
10 The Lewis Twist: Mind Over Matter	177

11	The Bare Necessities	193
12	Conclusion: Metaphysics as Intellectual Ergonomics	217
	<i>Appendix: Was There Anything Wrong with Psychologism?</i>	223
	<i>Notes</i>	226
	<i>References</i>	284
	<i>Index</i>	296

For Hilary Putnam

Acknowledgments

Many people have been kind enough to read and comment on drafts of the manuscript. Ben Crowe, Christoph Fehige, Leslie Francis, James Hawthorne, Erik Johnson, Kathrin Koslicki, John MacFarlane, Cei Maslen, and Paul Teller read through the whole thing. Among those who read one or more chapters—and my apologies if I’ve lost track of anyone—were Janet Abbate, Chrisoula Andreou, Jody Azzouni, Paul Benacerraf, Sarah Buss, Alice Clapman, Alice Crary, Kate Elgin, Don Garrett, Matt Haber, Eric Hutton, Rachana Kamtekar, Rosanna Keefe, Clifton McIntosh, Ron Mallon, John Martin, Matt Mosdell, Ram Neta, Lex Newman, Shaun Nichols, Bernhard Nickel, Laurie Paul, Anya Plutynski, Hilary Putnam, Diana Raffman, Michael Ranney, Guy Rohrbaugh, Sherri Roush, Laura Schroeter, Bill Talbott, Mariam Thalos, Valerie Tiberius, Johan van Benthem, Candace Vogler, and Gideon Yaffe.

For very helpful conversation, I’m grateful to Irene Appelbaum, Alyssa Bernstein, John Biro, David Chalmers, Charles Chihara, Heather Douglas, John Etchemendy, David Finkelstein, Peter Galison, Michael Glanzberg, Peter Godfrey-Smith, Bob Goodin, Peter Hylton, Jenann Ismael, Gabriel Richardson Lear, Dorothea Lotter, Chris Martin, Emily Rosenzweig, Steve Stich, Jim Tabery, and Jamie Tappenden.

Bits and pieces of this material were presented in various venues, including Kansas State University, the EMST Reasoning Group at UC/Berkeley, the University of Chicago, Vanderbilt University, the Center for Advanced Study in the Behavioral Sciences, Stanford University, the NYU/Columbia Vagueness Reading Group, the IUC Conference on Vagueness, the University of Miami, the University of Oklahoma, the University of Tennessee/Knoxville, Brigham Young University, King’s College London, and Reading University. I appreciate the audiences’ thoughtful feedback.

For drafting the illustrations, I'm grateful to Joe Ulatowski, and for editing assistance, to Erin Menut.

Work on this book was supported by fellowships from the National Endowment for the Humanities and the Center for Advanced Study in the Behavioral Sciences; I am grateful for the financial support provided through CASBS by the Andrew W. Mellon Foundation. I would also like to thank the Rockefeller Foundation's Bellagio Study and Conference Center, the University of Washington's Helen Riaboff Whiteley Center, and All Souls College for residential fellowships, as well as the University of Utah's Research Council and College of Humanities, for travel and sabbatical support.

Introduction

When you call something a hard truth, it's usually an announcement that you're about to say something your audience doesn't want to hear. That's not what the phrase is going to mean here: a hard truth is one that is just flat-out true (and if it weren't, it would be just flat-out false): not sort of true, not true in a way, not true up to a point, not somewhat true, and not any of the indefinitely many other variations on falling short of full truth. But hard truths in the more conversational sense are nonetheless in the offing, because hard truths, I will argue, are a lot fewer and farther between than philosophers have generally supposed, and to get them, you have to do things that philosophers have preferred not to think about or mention in polite company, things that you really can't do too much of the time. It follows that people have to do most of their thinking with the other kind of truth, namely, the truths that are only kind of true. The first part of this book will be given over to motivating the turn to partial truth; the second part will develop the argument for the claim I have just made.

There's a consequence to be drawn, to which the tail end of the book will be devoted: that a great deal of metaphysics is going to have to be substantially revised. That will be unwelcome news to professionals at one end of the philosophy business, but there will be a compensating payoff at the end: a new conception of what metaphysics *is*. By way of persuading you that the apparently not-so-sexy topic of partial truth is worth your attention, I'll start off this Introduction by explaining why that's a payoff to which we should look forward. I'll then quickly sketch some of the main ideas to be expressed by the core argument, and give a chapter-by-chapter overview of the proceedings. I'll wrap up the Introduction with a heads-up regarding the book's stylistic cues and organization.

1.1

Metaphysics (by which I mean not the bookstore shelving category, i.e., a euphemism for superstition, but rather one of the central and perennial concerns of philosophers) has gotten itself a reputation for being the most rarefied, abstruse, and impractical of pursuits. It is thought of not just as a philosophical specialty of merely intellectual interest, but one that does not even serve any intellectual interests other than its own: if you do not *happen* to be curious about such rarefied questions as, oh, what universals are, or what necessity is, or what causation is, or what objects are, and curious about them on their own account, you have no reason to care. Worse, this is the discipline's current self-conception, and not just a prejudice on the part of ignorant outsiders.¹ Metaphysics used to claim that it was the Queen of the Sciences, at a time when that was a proud claim to make, but it now presents itself as the superfluous royal figurehead of a Great Britain, Holland, or Sweden: on display, surrounded by much pomp, decorum, and ritual, and acknowledged to be absolutely useless.

I think that metaphysics is much more interesting and important than that, and if the argument I am going to be developing is correct enough, the conception of metaphysics I've just sketched is a mistake. Metaphysics, or most of it, should be thought of as an *applied* science, one part of which is concerned with determining what approximations or idealizations to use, and the other, an activity with affinities to product design. Metaphysics properly done (or, again, at least most of it: in the Conclusion, I'll return to the question of whether this is *all* of it) is an especially deep form of engineering, one that matters precisely because it is so useful as to be indispensable.

If you survey the history of philosophy, you will find that, throughout, metaphysics has addressed itself to the question: What does the world have to be like, if we are to be able to think about it? The question has come in for many substantively different readings over the course of its history—contrast the very different spins that, for instance, idealists and realists have put on it—but their shared content explains why metaphysics is not the silly enterprise of armchair physics it sometimes seems to be.² A philosopher's view of what thought consists in is a picture (or, in especially self-aware and ambitious cases, a theory) of rationality. So, not unreasonably, analytic metaphysics—metaphysics as it is done by analytic philosophers—has been largely a projection onto the world of a widely shared if not very clearly articulated picture of rationality.

In this picture, deductive inference, as first axiomatized by Frege and Russell, and as now taught in standard introductory logic classes, is first-class reasoning, and the prestigious core of rationality.³ Around it are clustered various second-class forms of reasoning, some still deductive, like modal logics and tense logics, and some that, while no longer extensions of the traditional deductive logic, are still formalized, like Bayesian inference, expected-utility theory, and so on. As we move farther away from the deductive core, we find still less prestigious modes of reasoning, such as inductive and abductive inference, and once we get far enough away, perhaps to arguments by analogy, the honorific title of reasoning is withdrawn. Deductive inference, and specifically the deductive logic taught in all those introductory classes, is *real* inference, and certainly the central, most important variety of it; the closer a form of inference is to such deductive inference, the more real it is.

That picture presupposes two further and related doctrines. First, that reasoning, properly performed, is a matter of eliciting true conclusions from true premises,⁴ and second, that truth is all-or-nothing: sentences and beliefs are either flat-out true, or flat-out false. Call the latter idea, that all truth is hard truth, the *bivalence doctrine*,⁵ call the former idea *validity as truth preservation*. Philosophers today tend to think of these as different *types* of claim: the one is about inference, and the other about semantics. In Chapter 3, I will explain why they travel together; for now, it is enough that the deductive logic taught in the standard logic class requires both that there be only two truth values, true and false, and that the notion of the validity of an argument is standardly introduced with some such gloss as: necessarily having a true conclusion if it has true premises.⁶ To anticipate, in domains of which either or both doctrines were false in the right sort of way, deductive logic would have limited or no applicability, and if most of our thinking were done in such domains, then, as far as practical importance goes, deductive logic would not capture the central or most important forms of inference. If analytic metaphysics really is a projection of this picture of rationality, then coming to see deductive inference as an unusual (though important) special case, rather than rationality's center of gravity, should be expected to have the effect of repositioning traditional views in analytic metaphysics as smallish regions of larger and more varied terrains.

These two doctrines—that valid inference is truth-preserving, and that truth is bivalent—are the philosophically standard view of these matters. But a second glance suggests that they cannot both be right; all too often, the starting points of our inferences are claims that we ourselves take to

be only approximately true, or technically true, or officially true, or true for present purposes. I am going to argue that we cannot do without such reasoning, that consequently truth must often be partial, that there must be a way to reason from false premises to conclusions that are, if not true, at any rate satisfactory, and so that we need to have a suitable notion of correctness of inference: that is, a notion of validity that does not amount to truth-preservingness. And I am going to argue further that when the hard (i.e., bivalent) truths of the title are available, special explanation is normally required, specifically, an engineering explanation. The perceived centrality of deductive logic, and the work in metaphysics that is tied to that perception, are the concomitants of a misguided way of thinking about truth. And so I think that revising the way we think about truth is the best way to fix these problems.

Let me now preview the ideas at work in the argument against the bivalence doctrine and validity as truth preservation. In doing so, I will help myself to a bit of philosophical shorthand; if you're not familiar enough with Kant and Nietzsche for the gestures at their views to be helpful, don't worry about it: the moves will be fully spelled out in the sequel.

1.2

Here is a Kantian thought. Suppose the world is a messy, surprising, and irregular place, and things as they are in themselves could be *any* which way. If those things were fully responsible for the inputs that go to make up your experience, then those inputs could also be just any way at all. But inputs that are just any way at all don't make up *experience*; if your cognitive inputs were arbitrarily varied and overwhelmingly multifarious, you would be intellectually overwhelmed by them, and you would, in the end, fail to have anything amounting to a mind. Somehow the raw input must be preprocessed before you see it: it must first be poured into molds. Since experience is the result of pouring raw inputs into molds, the molds have to be there *first*, and so the molds themselves can't be chosen or readjusted on the basis of experience.⁷

Here is a correction to the Kantian thought. Because the world could be just any way at all, pouring your proto-experience into a fixed set of molds will not do. Whatever won't fit into the molds will end up not being part of your mental life; that is, it will become invisible to you. But if the world

could be just any way at all, you can't know up front (in philosophers' Latin, *a priori*) what it's safe to disregard; to ignore you-know-not-what is to invite catastrophic consequences. And if you try to make everything fit—by *forcing* it in—you will end up breaking the molds. In the Kantian picture, the molds are much of the structure of your mind, so that would mean breaking your mind. Something has to give.

Let's agree that while you will have to pour whatever is going to make up your experience into molds, they don't always have to be the same ones. You can throw out old molds, and install new ones, if experience so warrants; in fact, you will have to, if you are going to maintain a unified mind. Since disposable molds are neither *a priori* nor necessary, they don't have to be Kant's categories, and we need to keep an open mind about what will replace them; perhaps plain, ordinary names, descriptions, and so on will do most of the time; perhaps special-purpose mathematical descriptions will sometimes be needed to exhibit them.

The corrected Kantian thought needs still further correction. If you could *always* tailor your repertoire of descriptions to what the world serves up—that is, choose molds to fit your experience—you wouldn't have to worry about things getting left out, or about smashing your mind on the world. If the world is as messy as experience suggests, however, too many molds would then be one-use throwaways, which would also leave you short a mind. And if the molds and their contents reflected too much of the messiness of the world, they wouldn't fit together inferentially—that is, in ways that allowed you to reason your way to important and warranted conclusions. Thinking requires thoughts that are repeatable, and which fit together inferentially. The world is such that fitting our thoughts to it trades off against fitting them to one another; the workable compromise is to be generous about what counts as fitting into a mold, and to accept less-than-perfect castings. Our literal vocabulary—or one of them—for saying how our opinions fit the world is our truth vocabulary. Truths that are entirely true—not true up to a point, not just sort of true, not only true in a manner of speaking, and not merely true *enough*—aren't suitable as contents for all, or even most of our thoughts, and we have to get along with soft, partial truth much of the time. This means that you can't always have hard truth.

One more round of revision to the thought we have been morphing. Sometimes you do need hard truth. Allow that Kant was right about this much: if unified intellects are to be possible, the world has to be poured into molds which guarantee that experience (or enough of it, anyway) will be intellectually tractable. If there are fairly close-in limits to how much

turnover we can handle in our supply of mental molds, and if we often have to be ungenerous about accepting an imperfect fit, then we will have to give up on the assumption that all the molds are inside or part of your mind. Experience must be shaped into the required configurations *outside* the mind. This idea is perhaps not entirely alien to the development of Kant's own thought. On his way to the central insights of the *Critique of Pure Reason*, Kant wrote that

if that in us which we call 'representation' were *active* with regard to the *object*, that is, if the object itself were created by the representation (as when divine cognitions are conceived as the archetypes of things), the conformity of these representations to their objects could be understood.⁸

Kant seems to have dismissed this option out of theological modesty; I take it he assumed that getting this solution to work would require omnipotence. And if all objects, all the time, had to be created in the course of representing them, the solution would in fact be unmanageable. But if we need to take this approach only some of the time, then engineering rather than omnipotence will often suffice.

To make thinking that deploys hard truths possible, we have to reengineer the world, by manufacturing objects to conform to our representations. When hard truth is uniformly available with regard to some subject matter, that's normally because we have constructed it. That last claim is easy to take the wrong way these days; but when I say that our truths are constructed, I don't just mean the wishy-washy kind of construction—'social construction'—that you hear so much about from postmodernists. I mean, in the first place, the kind of construction that involves power tools and earth-moving equipment. And so here is how our thrice-revised Kantian thought has come out: having a mind requires both thought that deploys partial truths, and altering some of the objects of thought to make full truth about them possible.

Here is a Nietzschean thought. There are far too many things in the world for us to have dedicated representations of each and every one. So if we are to be able to think about more than a tiny handful of them, we will have to use the same representation for more than one thing. Nietzsche held, at one point, that predication was really just this: the word 'leaf' begins its life as a name for a particular leaf, but is then applied to indefinitely many leaves.⁹ But because everything differs from everything else in indefinitely many ways, our predications are falsifications, or, as he moralistically put it,

'lies in the extra-moral sense.' When we call different items 'leaves,' Nietzsche thought, we are forgetting the multitudinous ways in which subsequent leaves differ from that original leaf.

Here is a correction to the Nietzschean thought, one that will strike any sensible contemporary philosopher as simple common sense. When we turn a name into a predicate, we selectively mask most of the detail (and not just the detail) in the original particular: when we call something a 'leaf,' we mean only to say that it resembles the original in precisely one respect, that of being a leaf. We institutionalize this kind of selective masking by introducing a new kind of item, the *property*, to be what it is that a predicate 'names.'¹⁰ Once the restraint in our semantic intentions is appreciated, predication is quite obviously not a falsification or 'extra-moral lie,' after all. What is left of Nietzsche's insistence that we arrive at predication by 'forgetting' is that the choices about which details to mask are very often made thoughtlessly, and consequently unintelligently; as we will see, this is not an unreasonable concern.

Let's run with the Nietzschean thought for a moment. There are too many properties in the world for us to have representations for all of them. If too much of the world is not going to be left invisible to us, we are going to have to recycle our predicative representations. Partial truth is the way (actually, one way) we do this: we apply a predicate when it does not quite apply, that is to say, if you insist on being picky about it, when it really does not apply at all. So-called partial truth is just another lie—to be sure, only in the extra-moral sense. But partial truth allows predication a penumbra in which numerous interstices of the world that would otherwise have remained in darkness are made visible to the intellect.

Here is an analogous correction to the amplified thought. When we advance a predication as a partial truth, we are selectively withdrawing our assent or commitment to the full reach or force of the initial predication. When we say that it's sort of true that it's a leaf, we intend to claim less than that it's a leaf—in just the way that when we claim that it's a leaf, we intend to claim less than that it's token-identical to some *ur*-leaf. Once the restraint in our semantic intentions is appreciated, partial truth becomes quite obviously not a falsification—an 'extra-moral lie'—after all. Partial truth is as innocuous a device as predication, and in a deep sense a very similar device: a means by which a mind whose representational capacities are necessarily quite limited can be cognizant of the plenitude of the world.¹¹ A philosopher who thinks that predication is fine but partial truth is not occupies an unmotivated and unsustainable position: the pressures

that take you from a language (or a ‘language’) composed of proper names for individuals to a language with terms for properties carry you past that position, to a richer language with expressive devices for partial truth. But we should not forget the analog of Nietzsche’s criticism of our choices of predication: that too often we choose our partial truths thoughtlessly and unintelligently.

1.3

Here is how I intend to develop and exploit the ideas I’ve just described. A substantial amount of additional warm-up is called for, and the first part of the book will be devoted to it. On the one hand, non-philosophers are unlikely to share the conception of truth that I will be contesting; they will wonder why we need an argument for the self-evident fact that many of the things that we think and say are only partially true, and will suspect that a fact as obvious as that is unlikely to have consequences which are both significant and unnoticed. Philosophers, on the other hand, will for the most part divide up into two groups: those who find bivalence self-evident, and specialists who take my concerns to have been already addressed by recent work on vagueness; both groups are likely to think that the claims I have just sketched are still far too impressionistic, and far too thin, to justify launching into an argument against bivalence. So I will begin, in Chapters 2 and 3, by laying out more of the big picture. I will try to make plausible the claim that bivalence is typically the product of engineering, and that because not everything can be engineered, we need to get along without bivalence a good deal of the time. Then I will introduce and criticize what is currently the most popular family of theories of truth (that is, theories that try to answer the question of what truth *is*).

Theories in the family are generically called ‘deflationist,’ and seem to imply that treatments of truth should not have metaphysical consequences. Since I mean to derive such consequences from my own treatment, I have to say why deflationist accounts of truth do not derail my arguments and my project more broadly conceived. I won’t exactly argue against those theories, but I will claim that the best stance to take toward them is diagnostic; once we place deflationist theories of truth in their historical context, their plausibility will diminish, and the associated idea, that there is not much you can *get* by thinking about truth, will be seen to be a mistake. More importantly,

by contrasting the function of truth, as deflationists understand it, with the functionality I am going to ascribe to partial truth, I will be explaining what I mean partial truth to be. (However, the reader should be warned that the official definition of partial truth will only be given in Chapter 6.)

Through this point, the focus will have been on motivating the view, rather than on arguing for it. The second part of the book takes up the latter task. I will make two complete passes over the book's core argument, in Chapters 4 and 5. The claim for which I will be arguing, that we reason using partial truths, belongs to logic in the pre-twentieth-century sense: it is a claim about how to think, or about the Laws of Thought. The analytic tradition has historically treated logic as a branch of mathematics, and for the most part let the prescriptive subject matter alone. Not surprisingly, the techniques for thinking about logic so understood have grown rusty with disuse, and one of the topics that has to be on the table is just *how* to conduct such an argument. Accordingly, as I lay out the argument, I will also be explaining what the form of the argument is, and why it is appropriate.

The core argument of the book will be (to appropriate a bit of an older vocabulary, which I will explain in due course) transcendental, and it is a remarkable feature of transcendental arguments that they can normally be put in either of two forms: as a task analysis (that is, the central claim in such an argument has the form: in order to do *this*, you also have to be doing *that*) or as an analysis of what is required for one or another aspect of unity of the self. Chapter 4 will develop the task-analysis version of the argument, to the effect that inference deploying partial truth is legitimate, because you can't make much headway in figuring out the facts without it. And it will also explain how the fact that you can't get along without it shows that an inference form is legitimate. Briefly, I will claim that the elements of thought must repeat themselves if thought is to be possible at all; the world is too messy for elements of thought to repeat themselves often enough, if they merely follow the world; if these two claims are correct, we must abandon both the bivalence doctrine and validity as truth preservation. In the course of the argument, I will also explain why the world has to be messy (and what is involved in my claim that it is).

In Chapter 5 I will throw the argument into the other form taken by transcendental deductions. I will first explain why the preconditions for (aspects of) unity of the self can determine what one's logic—one's guidelines for thinking—must be. I will introduce the notion of a unified intellect, and argue that the beliefs of a disunified intellect are not worth having. I will argue that minds of roughly our cognitive horsepower that do not reason

using partial truth fail to manage unified intellects: unified intellects are possible only if the bivalence doctrine and validity as truth preservation are false. Because the world is messy, fitting mental representations to the world does not allow them to be inferentially fitted to each other; partial truth, that is, allowing slack between our representations and the world, is what we need for the bits and pieces of an intellect to be glued together into an intellectually effective mind.

Chapter 6 will complete the core argument by turning to the philosopher's notion of belief. This discussion will allow me to provide, albeit belatedly, a definition of partial truth. It will allow me to address what I have found to be the most frequently pressed objection to the core argument. And it will provide, additionally and incidentally, the first of a series of applications, an illustration of my earlier suggestion that much analytic metaphysics is a projection of a theory of rationality: here, nearly exclusive attention to one kind of inference is responsible for a theory of the contents of people's minds.

Just to provide a sense of what I mean by that, let me pause to sketch how that application will go. Analytic philosophy of mind is focused very largely on beliefs. Belief-desire psychology is the dominant view, roughly, that all you need to talk about, when reconstructing the mental activity that amounts to reasoning, are beliefs and desires;¹² but desires do not have to do with truth at all, and are, on this kind of view, elements of practical rather than theoretical reasoning. So theoretical reasoning—that is, reasoning about what the facts are—is thought to be made up of beliefs, and beliefs alone.¹³

To believe something is to take it to be true, and so the assumption that there are only hard truths explains (much of) the contemporary philosophical commitment to belief psychology. If I am right, and there is a good deal that you only half-take to be true, then there must be steps in your reasoning that are not beliefs. Belief psychology (the theoretical side of belief-desire psychology) is appropriate for the special case of reasoning about hard truths. But if, as I am going to argue, hard truths are relatively rare, then beliefs are a good deal less common than is ordinarily believed; they turn out to be best understood as a special case of a broader class of mental items involved in theoretical reasoning. Thus the argument against belief psychology extracts a consequence for the metaphysics of mind from the previous chapters' theses about truth, and so bears out my suggestion that revising one's theory of truth promises important metaphysical payoffs.

With both the account of partial truth and the argument for it in place, I will be in a position to address the most prominent competing approach to the family of phenomena, and I will devote a chapter to recent work on vagueness. The literature is sufficiently specialized so that readers who are not already immersed in it may want to skip this third part of the book. Nonetheless, I will introduce the material I discuss in a way that I hope will make it accessible to those who are encountering it for the first time. And there is work that this chapter does over and above the necessary but unpleasant job of explaining why the apparent alternative is not really live: that of further filling in and supporting the picture of partial truth as coming in indefinitely many qualitatively different varieties.

The fourth and final part of the book pursues the point that there are philosophical consequences to these claims about truth and inference. Because there are many respects in which adopting them can change your substantive metaphysical position, for the first time in the course of the exposition there is a choice to be made as to which further topics to pursue. (Unlike the treatment of belief in Chapter 6, which is forced by the core argument, here I am selecting among many possible applications.) The approach I opt for is historical, and in Chapters 8–11, I will recount the development of one strand of twentieth-century analytic metaphysics, and exhibit how its turning points presupposed the bivalence doctrine. (The central figures in the narrative will be W. V. Quine, Donald Davidson, and David Lewis; that narrative itself, once embarked upon, will not allow much in the way of freedom of choice of topic, except as regards illustrations of its influence on other philosophers.) It will follow that giving up the bivalence doctrine means giving up the positions and methodologies constituting that strand. While the series of positions I will discuss was not by any means all of twentieth-century analytic metaphysics, it was structurally central, framing and motivating subsidiary metaphysical projects, many of which are still alive today, and bivalence figures into it in perhaps as straightforward a manner as we could ask. Abandoning that structurally important strand means abandoning the subsidiary projects as well. In short, I hope to establish that if we reject the bivalence doctrine and validity as truth preservation, then we will have to recognize a great deal of twentieth-century analytic metaphysics to have been a mistake.

If much of the metaphysics in circulation is an error, the natural next question is whether and how it can be replaced. So, finally, Chapter 12 will assess the outcome of the argument. I will return to the claim that metaphysics is a projection of a theory of rationality—a claim I will at that

point be in a position to cash out—and I will ask how metaphysics should be reconceived, if I am correct in thinking that hard truths are hard to come by. As already announced, I will argue that we should understand metaphysics as an applied science, one that can and should make our lives (our *day-to-day* lives) better than they now are.

1.4

Those familiar with the conventions of a literary genre allow their reading to be guided by its stylistic cues, but they do not normally pay explicit attention to the cues themselves. Because in this case the cues are likely to be misleading, let me take a moment or two to preempt a handful of misunderstandings.

First, in the business of professional philosophy, publication for the most part divides up into professional journals and books meant to be read solely by specialists, and popular or textbook expositions of material already presented in the professional format. Accessible writing (and especially, rehearsing background that a professional would not need) is accordingly taken to signal that most of argumentation is supposed to have happened elsewhere. Put a bit more bluntly, you're not supposed to take philosophical prose seriously unless nobody who wasn't already a philosopher of the same tribe could read it comfortably.

My own view is that philosophy should always be written to be as broadly accessible as possible, and in this case there is a special reason for making the writing readable to people who are not versed in the current philosophical subspecialties. One intended upshot of the argument of the book is that the approaches taken in a number of philosophical specializations—theory of truth, philosophy of logic, metaphysics, and, incidentally, epistemology—are deeply misguided. If that very ambitious follow-on thesis is taken to heart, philosophers will change the way they work in those areas. (I know, it's a long shot.) If they do, and if the book is written so that only professionals working in those fields, as they are now constituted, can read it, then if the book is successful, it will thereby make itself unreadable. Perhaps it's too ambitious to write for the ages, but I would at least like to be writing for a number of years down the road, even if (especially if) the argument is found convincing. That means writing readably, even to those who do not

have the background of contemporary specialists. (There are unavoidable exceptions; I have encapsulated them in notes and an appendix.)

To write for nonspecialists amounts to writing prose that undergraduates can read, and to a professional nowadays, that stylistic choice is likely to convey the impression that the discussion is underargued, and probably that it is loosely organized, as well. Please resist the impression; while the argument proper will not kick in until Chapter 4, it is meant to be as tightly controlled as it gets.

Second, because the argument is ambitious, both as regards its methodology and its conclusions, there are inevitably going to be a lot of pieces on the board. I have found that it pays to disabuse readers of the impression that I am meandering from optional topic to optional topic. Because the argument determines the organization of the material, prior to the turn to applications in Chapter 8, there are almost no open choice points. (A discussion of epistemology in Chapter 6 is an exception, but is there because almost all the work needed to get it going is required by the argument proper.) I will almost always explain what one discussion or another is doing where it is, but I will not always make a big deal out of it. Once again, I am trying to write accessibly, which means departing from the overly verbose conventions of the journal article.

Third and finally, there is an obstacle to following the flow of argumentation that is not merely stylistic. This book argues for a view about what arguments—theoretical arguments, viz., arguments as to what the facts are—have to be like. Professional philosophers are trained to argue from putative flat-out truths to further flat-out truths, and to make their arguments as close to deductively valid as they can. I am claiming that many arguments have to go from ‘true enough’ to ‘true enough,’ that those arguments will not be deductive, and that they have to be understood and assessed very differently from deductive arguments.

To the extent that the argumentation in this book is about matters of fact (though bear in mind that there are going to be a number of practical arguments, that is, arguments whose conclusion is about what to do), it will be almost entirely of the very type I am attempting to characterize and defend. That will make it harder for philosophers used to deductively cast argument to see it as argued (and to see it as *tightly* argued). I assure you that the argument is as tight as I can make it, but what it is to make an argument tight, when the argument turns on partial truths, may be different than what you are used to.

Part I
Motivation

The Truth in Bivalence

I've suggested that many philosophers have grown up with the bivalence doctrine: the idea that what we assert—internally, in our thoughts, or externally, in our writings and utterances—is either just plain true or just plain false. I think it's fair to say that, with the exception of a few groups of specialists, the bivalence doctrine is very widely held among philosophers, and most of the time, just taken for granted.¹

However, it's not nearly as widely held among non-philosophers (and philosophers when they're not wearing their philosopher hats). Almost anyone will balk at answering a good many yes-or-no questions with a straight yes or no. Do I want to go hiking tomorrow? *We-ell*, it's hard to give a straight answer to that question. Is the defendant guilty? A lot of jurors would *really* like the option of saying, 'Sort of...' When I am told that Nashville is four hours from Memphis, I take that to be more or less true, but not exactly true. And informed that the dog ate the homework, I may understand that to be false in spirit but not in letter (if, for instance, I know that the student has fed the unfinished homework to the dog); that is, I may take it to be true...in a way. Just going by what ordinary folks say, the bivalence doctrine is false. So who's right, the philosophers or the laity? Is the bivalence doctrine true?

In this chapter, I'm going to suggest that, even though it's not true in the main, the bivalence doctrine has some truth in it. I propose to take a shot at teasing out what it is the bivalence doctrine gets right; on the way, I will try to show how it goes wrong, and to say why it's nonetheless important to appreciate the truth in bivalence.

2.1

First, terminological preliminaries. I'll use the term 'bivalent' in two related ways. Of sentences or statements or thoughts or assertions, taken one by one, I will say that they are bivalent if they are guaranteed to come out flatly true or flatly false when they are held up against some range of backgrounds in which we realistically expect them to be used. (E.g., the background to my use of sentences containing the word 'bald' is a world in which some people have no hair, some have just a little, some have more, and some have very full heads of hair.)² Now, although in the course of dealing with our day-to-day affairs, we're usually interested in the truth and falsity of sentences or thoughts, taken one at a time against the actual circumstances in which they present themselves, here we will be concerned primarily with the counterfactual or dispositional behavior of categories of sentences. That is, we'll want to know, not whether the sentence 'John Smith is bald' is true, or even whether it is sure to come out true, but whether, for instance, the class of sentences containing the predicate '... is bald' can be counted on to come out flat true or flat false.³ So, and here's the second usage, I will say of categories of sentences or thoughts, that they are bivalent, when each of the sentences in the category is bivalent.

Let me mention four consequences of using our terms this way: First, because bivalence is, here, a dispositional concept, a sentence can fail to be bivalent even if, in a particular case, it comes out flat true or flat false (in just the way that a glass can be fragile even while it is still unbroken). Second, and closely connected to the previous point, a category of sentences that is not bivalent can contain sentences that are bivalent, taken individually: the category of sentences constructed using the predicate '...is bald' is not bivalent, even though 'Somebody or other is bald' is, realistically, inevitably true, and so bivalent. Third, bivalence so understood is sensitive to what is being taken to be the relevant background: in a world somewhat different from our own, in which in-between scalps did not occur, the category of sentences formed using the predicate '...is bald' would be (as far as the predicate went) bivalent after all. And fourth, bivalence, treated this way, can be thought of as a matter of degree, or something like a matter of degree; if *almost* all the occasions for the use of some sentence make it come out fully true or fully false, we may want to ascribe, say, almost complete bivalence.

2.2

We can proceed by treating the layman's view as a working hypothesis. Now, taking seriously the hypothesis that sentences, thoughts, and so on often fail, in one way or another, to be bivalent, means allowing that there is room for explanation when categories of sentences or thoughts *are* bivalent. If bivalence should not be taken for granted, then when one has a thought that is in a position to be either true or false, we can ask how that came about. One way of taking this claim makes it uncontroversial. It is a familiar fact that to have a thought that is true, one must first have a thought that is clear enough to be true or false—a thought capable of engaging the facts that would make it true—and we know from experience that attaining the requisite clarity is hard work. If you're a philosopher by trade, think of the struggle involved in converting your raw and inchoate inspiration into a claim with enough definition to stand a chance at being true or false. And if you're a philosopher with an academic day job, think of the papers your students submit: too often, the problem with student papers is not that they are *wrong*, but that they haven't gotten to the point where they might be so much as wrong. It's no accident that philosophers often think of the alternative to full-truth-or-full-falsity as vagueness.⁴

Sometimes, getting into the position of having a bivalent thought is substantially or entirely a matter of arriving at a clearer formulation of it. And there is a widely held view about beliefs which suggests that this is how bivalence is always managed. On the view in question, beliefs are to be understood in terms of their 'direction of fit': beliefs are supposed to fit the world, and are contrasted with desires, to which the world is to be made to conform. In other words, when there is a mismatch between the world and a desire, as when you want your pantry to be well-stocked, but it is not, the mismatch is addressed by getting provisions to put in the pantry; whereas if you believe your pantry is well-stocked, and it is not, the mismatch is properly addressed by changing your mind as to the pantry's contents.⁵

Consider a case where the mismatch between belief or statement and world amounts to a failure of bivalence. For example, I point out that the cloud up there is shaped like a horse's head. This is not, strictly speaking, true; no horse's head was ever shaped *quite* like that.⁶ It is true enough, true in a way, true for present purposes. Now, if for some reason I were to decide that this kind of half-fledged truth was not good enough, an understanding of belief in terms of direction of fit would likely prompt me to address the