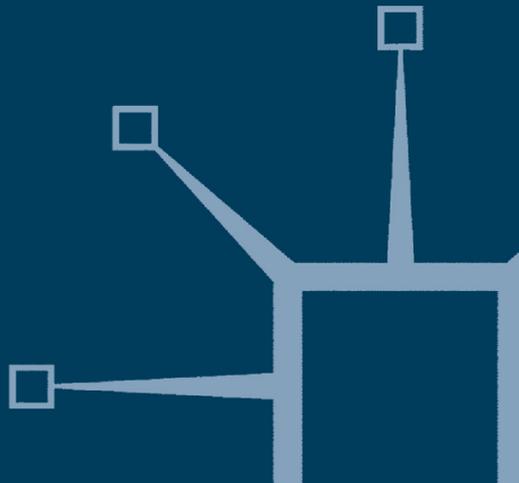


palgrave
macmillan

Aiming at Truth

Nicholas Unwin



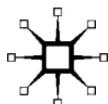
Aiming at Truth

This page intentionally left blank

Aiming at Truth

Nicholas Unwin

palgrave
macmillan



© Nicholas Unwin 2007

Softcover reprint of the hardcover 1st edition 2007 978-0-230-50683-1

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No paragraph of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright, Designs and Patents Act 1988, or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, 90 Tottenham Court Road, London W1T 4LP.

Any person who does any unauthorised act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

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

First published 2007 by
PALGRAVE MACMILLAN

Houndmills, Basingstoke, Hampshire RG21 6XS and
175 Fifth Avenue, New York, N.Y. 10010

Companies and representatives throughout the world

PALGRAVE MACMILLAN is the global academic imprint of the Palgrave Macmillan division of St. Martin's Press, LLC and of Palgrave Macmillan Ltd. Macmillan® is a registered trademark in the United States, United Kingdom and other countries. Palgrave is a registered trademark in the European Union and other countries.

ISBN 978-1-349-35307-1

ISBN 978-0-230-62490-0 (eBook)

DOI 10.1057/9780230624900

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources.

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

Library of Congress Cataloging-in-Publication Data

Unwin, Nicholas, 1954–

Aiming at truth/Nicholas Unwin.

p. cm.

Includes bibliographical references and index.

1. Truth. I. Title.

BD171.U59 2007

121—dc22

2006050350

10 9 8 7 6 5 4 3 2 1
16 15 14 13 12 11 10 09 08 07

To my mother and father

This page intentionally left blank

Contents

<i>Preface</i>	ix
Introduction	1
1 Truth and the Norms of Assertion	8
1.1 Why the nature of truth itself is irrelevant	8
1.2 An argument against truth-aiming	21
1.3 A revised argument	26
1.4 Intensional and extensional senses of 'aiming'	31
2 Scepticism – and How Not to Avoid It	38
2.1 Types of scepticism	38
2.2 The problem of induction	42
2.3 The underdetermination of theories by data	50
2.4 Externalist criticisms	63
2.5 Transcendental arguments	68
3 Cognitive Contingency	74
3.1 How to be an alien	74
3.2 Alternative classifications	78
3.3 Alternative conceptual schemes	85
3.4 Radical interpretation	91
3.5 Further objections	98
4 Belief and Acceptance	108
4.1 Acceptance as a distinct attitude	108
4.2 Voluntariness and 'directions of fit'	117
4.3 Degrees of belief	130
5 Belief and Knowledge	138
5.1 What is knowledge?	138
5.2 An 'ecological model'	144
5.3 Justification regained	155

6	Pragmatism and the Value of Truth	162
6.1	Rorty on pragmatism	162
6.2	The autonomy of belief	169
6.3	Wishful thinking	177
7	Relativism and the Limits of Conversability	185
7.1	Relativism and error	185
7.2	Tolerance	191
7.3	The individuation of cultures	197
8	Mind, World and Realism	203
8.1	Types of realism	203
8.2	Transcendental idealism and internal realism	205
8.3	Representation and the 'Olympian standpoint'	217
8.4	Conclusion	225
	<i>Notes</i>	227
	<i>Bibliography</i>	237
	<i>Index</i>	244

Preface

The ideas in this monograph were developed over many years. Earlier versions were published as 'Beyond Truth: Towards a New Conception of Knowledge and Communication', *Mind* 96 (1987), pp. 299–317; and 'What Does It Mean to Aim at Truth?', *American Philosophical Quarterly* 40 (2003), pp. 91–104, a much expanded version of which appears as Chapter 1. I am grateful to both journals for granting permission to reprint.

I am indebted to many people over the years who have provided encouragement and assistance. These include several anonymous referees who have given detailed and helpful reports on the above articles as well as earlier versions of the whole book. They also include people who attended the papers which I delivered, at various times, at seminars at the Universities of Bolton, Bradford, Glasgow, Hanover, Lancaster, Manchester, Manchester Metropolitan, Nottingham, Sheffield and Southampton. Alex Bird and Richard Holton were especially helpful. Leslie Stevenson wrote a useful and perceptive commentary on my 1987 piece. I was fortunate to be taught Kant at Oxford by Ralph Walker, and his version of transcendental idealism played an important role in stimulating many of the ideas in this book. However, most particular thanks are due to Julian Dodd, who read through several drafts of the manuscript, and made many useful suggestions as well as some very encouraging comments. I am also grateful to the University of Bolton for teaching remission in recent years during which this book was developed and completed.

University of Bolton
May 2006

Introduction

Most of us feel that truth is something important, even if we cannot easily say why. If someone believes something which we think is not true, then our instinct is to try to get this person to change his or her mind. Likewise, if we suspect that our own beliefs are not true, then our instinct is to try to replace them with beliefs which are. Of course, we all know people who seem to care very little about these matters, who routinely confuse genuinely believing that something is the case with merely hoping that it is, and who are too stupid or lazy to try to reform their various prejudices; but we typically suppose that we ought not to do likewise. Rather, we should aim at truth.

Yet explaining why we should do this, and what is involved in doing so, proves to be a surprisingly difficult undertaking. For one thing, it is debatable just where the main issues reside. Is it in the nature of truth itself? Or in the nature of belief? Or is it a mistake to suppose that truth as such is what matters, and that we should instead concentrate on related concepts such as honesty and intellectual curiosity? It may seem obvious that true beliefs are better than false ones, but some have argued that, although the natural sciences are generally thought to provide the most reliable methods of finding things out, scientific theories do not really aim at truth at all, and that they merely help us to predict and control our environment. Furthermore, it is often claimed that this is all that we could ever sensibly want them to do. If this idea is generalized, then a number of peculiar options suggest themselves. It may be thought, however, that such considerations can only be of marginal interest, and that the desirability of believing only what is true remains justified in at least most ordinary instances. This, indeed, is the conventional view.

Yet there are deep theoretical difficulties here. Even if we were to grant that truth should be aimed at (and that we know what it means

2 *Aiming at Truth*

to do so), there are many old and persistent arguments that appear to show that this task is impossible to fulfil. These are the arguments of the sceptics. From ancient times, it has been held that we cannot, even in principle, tell whether our beliefs are really true, and this has led some to conclude that we should abandon them altogether. This may sound utterly foolish, but the arguments for scepticism are powerful, and it is far from clear how they are supposed to be met.

In this book, we shall be concerned with a particular type of sceptical argument, which runs as follows. We are finite intelligences who obtain such knowledge as we have about the world, primarily through our senses. This means that the actual information that we receive from outside is very limited. For example, our visual information is based merely on stimulation of the retina. Since the world is three-dimensional, and our retinas are only two-dimensional, a massive amount of data is inevitably lost. Somehow, despite this, our brains manage to 'fill in the gaps', and we are gradually learning more about how they do this from work done in psychology, neurophysiology and elsewhere. However, a philosophical problem suggests itself at this point. Suppose that we had had brains of a very different kind: would it not be likely that they would have filled in the gaps very differently? Perhaps there are intelligent creatures somewhere else in the universe whose brains do precisely that. If so, why should it be supposed that our method of filling in the gaps is better than theirs? More specifically, why should we suppose that our method of interpreting data is more likely than theirs to lead to the truth – that is match the original reality? The symmetry of the situation suggests that no good reason can be given, and this is clearly a cause for concern if we think that we should aim at truth.

Some conclude from arguments of this kind that the world we take ourselves to have knowledge of is itself some kind of subjective construction. This leads to the idea that, whereas we live in one such construction, these other creatures (should there be any) would live in another, and there is no single, all-embracing objective reality – and therefore no genuine conflict to resolve. Yet this strategy is not altogether satisfactory, if only because it leaves it utterly mysterious how these different subjective constructions are supposed to relate to each other. They must have some relationship, for these other intelligent creatures are presumably part of our own reality, or else they could never have entered the story – *our* story – in the first place.

This book is centred around an alternative strategy for dealing with this problem. On our view, it is not the notion of an objective reality

that needs adjustment, but rather the notion of a belief. Believing something is typically understood as demanding agreement (or, at least, non-disagreement) from everyone else, but if 'everyone else' includes these alien intelligences, then trouble instantly emerges. Our solution is to revise this ordinary notion in such a way that it no longer demands such widespread agreement. Rather, we introduce a perspectival element into the heart of our cognition, and in an unusual way. The *contents* of our 'beliefs' remain quite unchanged, but the type of attitude itself that we have towards these contents is subtly altered. We shall argue that this strategy enables us to deal effectively with a number of central philosophical problems. These include not only the problems about scepticism and objectivity that we have just mentioned, but also issues that are more traditionally associated with the thesis that beliefs aim at truth, such as whether one can believe at will, and the distinction between theoretical and practical reasoning (i.e. reasoning about what to think, as opposed to reasoning about what to do).

But will our 'beliefs' still aim at truth if they are modified in this way? The chief goal of Chapter 1 is to show that they will, but in a way that requires us to examine very carefully just what is meant by 'aiming' at something. The notion of an aim has both 'intensional' and 'extensional' senses, to use the traditional vocabulary,¹ and it is argued that alternative intelligences could still be aiming at truth in the former sense (though not the latter) even though they are, so to speak, aiming in a different direction to us. Moreover, this is not, oddly enough, because there is something mysterious about the concept of truth itself. On the contrary, we shall also show that this concept is utterly straightforward (or, at least, that if there are any problems attached to it, they have nothing to do with the matter in hand). Much traditional philosophizing about the nature of Truth and Reality will be argued to be misplaced, and we shall see that what is really interesting lies elsewhere.

Clearly, these ideas will not amount to much, unless it really is possible for there to be alternative intelligences with belief-systems which are genuinely incompatible with, and yet (in some sense) equally good as our own. The purpose of the next two chapters is to show that this possibility is genuine. Chapter 2 elaborates and defends the general argument for scepticism sketched above. Our chief focus here is on what Quine calls the 'underdetermination of theories by data' (roughly, the thesis that there is always more than one theory consistent with any given body of evidence), which in turn is a special case of the traditional problem of induction. This problem, which originated with Hume, challenges our right to draw any inferences that go beyond our

immediate experience. We can, of course, give a straightforward pragmatic defence of our right to adopt an inductivist strategy, for nothing else could conceivably work for us, but we argue that this is not enough to justify actual *belief* that such a strategy will work. This negative result is justified in part by developing what Goodman calls the ‘new riddle of induction’. This centres around the point that even if we do accept that inductive inferences are generally reliable, what they yield will depend on how the evidence is described. If we are prepared to describe it oddly – for example, by using Goodman’s famous (or infamous!) words ‘grue’ and ‘bleen’, instead of the orthodox ‘green’ and ‘blue’ – then we can apparently justify odd conclusions.² The reason why we find some descriptions odd and not others seems to depend ultimately on human psychology, and this leaves open the possibility that intelligent creatures with very different psychologies will thereby react very differently to a given body of data – and with an equal level of justification. This divergence could continue even if investigations carry on indefinitely, for beliefs about the world will always go beyond what is yielded by the direct evidence of sensory experience.

However, even if we agree that the world might, in fact, be very different from the way that it appears to us to be, this will not yield an interesting scepticism unless it could be shown that alternative belief-systems are a genuine option. We agree that they are not a genuine option for *us*, but our thesis depends on the possibility of alternatively intelligent creatures, whom we variously call ‘Martians’, ‘grue-users’ or (more neutrally) ‘aliens’. The burden of Chapter 3 is to show that we really can make sense of this idea. To succeed here, it is not enough merely to show that such aliens could, in principle, exist – though even that much is not easy, since they need to be very alien indeed in order to qualify for the role. In order for our account to have any real credibility, we must be able to show how it could, in certain circumstances, be reasonable for us to identify such an alien *as* an alien. This requires us to examine a number of influential arguments about the nature of language, understanding and interpretation, due to Davidson and others, which insist that such circumstances are impossible – even in principle. We shall argue, by contrast, that we cannot fully understand ourselves until we can locate our own particular intelligence within a wider space of alternatives, and that this requires a fundamental revision of many of our ordinary assumptions.

Having established that there is a serious sceptical case to answer, we need to show that our strategy manages to answer it. In Chapter 4, we look much more carefully at the concept of belief, and compare it

with the much weaker notion of acceptance. The former seems to differ from the latter in that it aims at truth, does not depend on context, and is involuntary. The weakened or 'reduced' belief that we wish to construct needs to be somewhere in between the two, and the chief task of this chapter is to show that we can add context-independence and involuntariness to acceptance without implying anything about truth. It is here that we discuss most fully what it is that makes it impossible to believe at will, and the question of whether truth-directedness reflects an intrinsic functional quality of belief-states or simply the aims of the believer. We also compare our sort of reduced beliefs with the 'degrees of belief' that are central to philosophical analyses of the concept of probability.

In Chapter 5, we attempt to show that our modified concept of belief is adequate for ordinary purposes. We examine more carefully the concept of knowledge, and show that the standard view of it, as something that logically requires the truth of what is known, is, surprisingly enough, incoherent. This is because knowledge, as ordinarily construed, is required to combine truth and optimal cognition in a unified and non-arbitrary manner; but our story about aliens ensures that this is impossible. We develop an alternative, 'ecological' model which emphasizes the importance of biological harmony between organism and environment, as opposed to a metaphysically guaranteed 'fit' between representations and what is represented. It is here that the most radical and revisionist aspects of our thesis become apparent. Nevertheless, we are able to show that our revised model works very well for us, and that its strange pluralist aspects – which are, in fact, very attractive in many ways – do not damage our day-to-day practices. Moreover, most of our ordinary 'beliefs' (when suitably re-fashioned) can now be seen to be justified, and the sceptic has been answered in a highly effective way.

Yet even if we agree that we can get past our sceptical problems, and thus 'aim at truth' in some suitable sense, the question still remains as to why we should do so. A certain kind of pragmatist might argue that we should be better off if we lost interest in truth altogether, and replaced all our beliefs with mere acceptances throughout. Such a policy may look unappealing, but we need to explain what, if anything, is actually wrong with it. This is the task of Chapter 6, which looks at certain types of connection between beliefs, acceptances and desires. We examine here critically the views of some contemporary pragmatists, notably Richard Rorty. It is also here that we consider what would happen if we were to blur the distinction between theoretical and practical reasoning.

It is concluded that it is surprisingly difficult to give a non-circular justification of the value of truth. For example, we argue that we should not assume that wishful thinking automatically involves self-deception, or that the value of truth derives simply from the value of honesty. And even if it is hard to imagine human beings who had no interest in truth whatsoever, it is unclear that this proves anything fundamental about truth; rather, it may just prove something fundamental about human beings.

In so far as we claim that there are many different, but equally valid, cognitive systems, we are clearly endorsing a version of relativism, a theory that is usually abominated by philosophers. The purpose of Chapter 7 is to show that our version does not suffer from the faults that give rise to this hostility. Specifically, we argue that a sharp distinction can still be maintained between what we actually think and say and what we ought to think and say, that criticism both of ourselves and of aliens is permissible (indeed, mandatory), and that, although this does lead to problems about the limits of toleration, this is a problem which is not peculiar to our theory. We also look more carefully at how sharply we can draw the distinction between one sort of intelligence and another, and conclude that there is a considerable indeterminacy here. However, we also argue that the lack of precise boundaries between ideal systems of thought need not present us with any major difficulties. Rather, it allows for a welcome degree of flexibility within human nature itself. 'Conversability' is Rorty's term for the willingness to engage others in conversation, including those with a very different outlook from our own. The sort of aliens that we are concerned with will evidently stretch the rules of conversation considerably, yet conversational rules prove to be decisive in explaining the difference between full belief and its weaker counterparts. We examine some of the implications here.

As mentioned above, we distance ourselves from theories that attempt to subjectivize reality, or to treat the ordinary world around us as some kind of mental construction. The philosopher whom we look at most carefully here is Kant. Although we consider his views on several occasions throughout the book, our main treatment is reserved for the final chapter. Here, we argue that our perspectivalist approach is, in fact, quite consistent with the commonsensical 'realist' view that the world is the same for everybody and completely independent of how anyone thinks about it. We also argue that our thought-experiments about Martians and so forth are of vital importance in clarifying the nature of representation and the relationship between mind and world. Our overall views

are located within a more general metaphysical discussion which also examines the views of several contemporary thinkers.

Issues about scepticism, objectivity and alien systems of thought play a dominant role in the argument throughout, and the central aim of this book is to defend a specific, though wide-ranging, metaphysical and epistemological thesis. The nature and significance of truth-aiming is discussed primarily within this context, and this is perhaps an unusual strategy. However, we shall see that it has considerable merits. The most obvious drawback with our proposed new world-view is that it appears, at first sight, to imply that aiming at truth is neither possible nor even particularly desirable. The traditional problems therefore take on a particularly intense and urgent quality when they are examined from within our framework. Furthermore, we shall see that, even when we examine what it is to aim at truth from a more conventional point of view, the problems that emerge invariably turn out to be closely related to the metaphysical and epistemological issues that motivate our own position. Topics that are ostensibly rather disparate thus turn out, in fact, to be intimately connected, and our approach shows that this is so in a particularly effective way.

All this may sound surprising. Yet it should be clear that the problem of scepticism is of some basic relevance here. Aiming at truth can be difficult, and it can sometime seem too difficult to be worth bothering with. It is striking that, although we know far more than ever before, much of present-day thought is characterized by a general lack of confidence in our right to assert our own opinions and to say that people with whom we disagree are wrong. Such loss of assertiveness can be a welcome antidote to dogmatism, but can also paralyse our critical faculties. Much of the hostility directed to what is often called 'postmodernism' is centred around these issues. There is a certain kind of all-pervasive scepticism that is very apparent nowadays. This book addresses some central problems which underlie these concerns. Scepticism of all kinds has deep metaphysical causes, and once we understand them, we can see our way to dealing with a number of pressing concerns that affect us all.

1

Truth and the Norms of Assertion

Before we can discover the extent to which it is either possible or desirable to ‘aim at truth’, we must first work out what it means to do so. This will prove to be a surprisingly difficult task. It might be thought that this is because the concept of truth itself is elusive but, oddly enough, the real reasons turn out to be quite different. Rather, it is the concepts of assertion and belief that underlie the problem here. The distinction between belief and acceptance is introduced, and we expound a particular kind of sceptical argument that seems to undermine the desirability of having any beliefs at all. Our solution, that our beliefs may be legitimized if they are ‘reduced’ in a special sort of way, is defended, and we outline the general metaphysical significance of our proposals.¹

1.1 Why the nature of truth itself is irrelevant

Philosophical debates about truth take many forms, but they tend to focus on two main issues. The first, and most discussed, concerns its nature; the second concerns its importance. All parties agree that truth satisfies certain very minimal conditions, such as the ‘Equivalence Schema’:

(ES) The proposition that p is true if and only if p .

This says, roughly speaking, that calling a given proposition true is equivalent to just asserting that same proposition. For example, $\langle\langle$ Snow is white \rangle is true \rangle is logically equivalent to \langle Snow is white \rangle : there is no possible situation in which one of these propositions could be true but not the other.² The chief problem about the nature of truth is whether ES says all that there is to say on the matter.³ Deflationists,

such as Paul Horwich (1990), maintain that it does. Substantivists, by contrast, maintain that truth is an important semantic, metaphysical or normative notion whose nature has yet to be fully explicated. The correspondence, coherence and pragmatist theories of truth are examples of substantivist theories. The chief problem about its importance concerns why we should aim at truth, why true beliefs and assertions are better than false ones.⁴ The problems are related in so far as it is a challenge to deflationists to explain how an utterly minimal notion can acquire the importance that it apparently has, though it is also clear that each problem raises issues that the other should not be expected to address. What is less often thought to be a problem, however, is just what aiming at truth itself actually amounts to. Nevertheless, it will be argued in this chapter that the matter is much less obvious than it seems, and furthermore, that it raises issues of great importance in the theory of knowledge, particularly concerning the development of effective anti-sceptical strategies.

The overall dispute between deflationists and substantivists is not our problem, and there may indeed be problems about the nature of truth that are not captured by ES. However, it is important to see that Horwich is at least right to think that any such problems are of no relevance to the question of why truth should be aimed at. The reason why it seems otherwise is that ES itself certainly has no normative implications, from which it apparently follows that truth must obtain its desirable features from elsewhere. However, this is an illusion. We may agree – at least for the moment – that we should indeed aim at truth. But what does this mean? The most natural analysis is along the lines of:

(TN) We should assert that p only if $\langle p \rangle$ is true.

We say ‘only if’ rather than ‘if and only if’, since the latter demands, absurdly, that we take an interest in everything. We shall see that this principle is far more controversial than it may appear; but let us disregard, for the moment, the question of whether it is actually true, and instead merely examine it as a proposition. It is certainly normative, and it is certainly about truth, but does it thereby expound a norm *of* truth? This does not follow, for although it is about truth, it is also about assertion, and it tells us something about how and when we should assert things. Once we understand it in this way, the fact that TN, our ‘truth norm’, cannot be derived from ES in no way shows the inadequacy of the latter, for ES was never supposed to tell us anything about assertion. Obviously, the concepts of truth and assertion are importantly linked,

but this does not imply that we cannot grasp the former until we have grasped the latter. It might just as easily be the other way round.

Horwich argues along similar lines, but he supposes that the primary reason that assertions have the normative features that they do derives from social and ethical considerations – the fact that we ought not to lie, for example.⁵ But such ethical issues will not arise unless assertion already has appropriately truth-directed features, and we shall see that much more needs to be said at an earlier point of the analysis.

What are assertions? What is it to invest a speech-act with assertoric force? We should note, first, that not every indicative utterance qualifies as an assertion. An actor on stage may utter a declarative sentence, but he does not *assert* the proposition that it expresses. Why is this? A simple answer is that assertions express beliefs, and theatrical utterances do not. Beliefs are meant to be true, whereas plays are meant to be fiction. This fact is independent of, and prior to, the social and ethical considerations mentioned by Horwich; it ensures that the ethics of lying has no relevance to what actors might say when on stage, for example. However, as it stands, this answer does not get us very far, since it merely transfers the problem to one of belief. To say that beliefs are *meant* to be true, or that they *aim* at truth, is doubtlessly correct in some sense, but we still have to explain what it means.

In fact, the two questions run together. Assertions are the linguistic expressions of beliefs, and the relevant norms governing them are essentially the same: we should believe that *p* only if $\langle p \rangle$ is true. However, it is often better to focus on the latter rather than the former, since it is easier to see what they might be contrasted with: we have a richer and more precise vocabulary here. I may *believe* that *p*, but I may also *imagine* that *p*, or *consider* (for the sake of argument) that *p*. More relevantly, many thinkers have argued that *acceptance* is something significantly different from belief itself, and with different normative constraints.⁶ For example, Michael Bratman writes,

Belief has four characteristic features: (a) it is . . . context-independent; (b) it aims at the truth of what is believed; (c) it is not normally in our direct voluntary control; and (d) it is subject to an ideal of agglomeration. In contrast, what one accepts/takes for granted (a) can reasonably vary . . . across contexts; (b) can be influenced by practical considerations that are not themselves evidence for the truth of what is accepted; (c) can be subject to our direct voluntary control; and (d) is not subject to the same ideal of agglomeration across contexts. So acceptance in a context is not belief.⁷