

# metaepistemology and relativism

j. adam carter

PALGRAVE INNOVATIONS  
IN PHILOSOPHY



# Metaepistemology and Relativism

*Palgrave Innovations in Philosophy*

Series Editors: **Vincent F. Hendricks**, University of Copenhagen and  
**Duncan Pritchard**, University of Edinburgh.

*Titles include:*

Mikkel Gerken

EPISTEMIC REASONING AND THE MENTAL

Kevin Meeker

HUME'S RADICAL SCEPTICISM AND THE FATE OF NATURALIZED  
EPISTEMOLOGY

Ted Poston

REASON AND EXPLANATION: A DEFENSE OF EXPLANATORY COHERENTISM

Aidan McGlynn

KNOWLEDGE FIRST?

E.J. Coffman

LUCK: ITS NATURE AND SIGNIFICANCE FOR HUMAN KNOWLEDGE AND  
AGENCY

Jonathan Matheson

THE EPISTEMIC SIGNIFICANCE OF DISAGREEMENT

Annalisa Coliva

EXTENDED RATIONALITY: A HINGE EPISTEMOLOGY

Rachel McKinnon

THE NORMS OF ASSERTION: TRUTH, LIES AND WARRANT

*Forthcoming titles:*

Annalisa Coliva

THE VARIETIES OF SELF-KNOWLEDGE

Julian Kiverstein

THE SIGNIFICANCE OF PHENOMENOLOGY

David Pedersen

POLITICAL EPISTEMOLOGY: EPISTEMIC THEORIES AND KNOWLEDGE  
INSTITUTIONS

Christopher Pincock and Sandra Lapointe (*editors*)

INNOVATIONS IN THE HISTORY OF ANALYTICAL PHILOSOPHY

---

**Palgrave Innovations in Philosophy**

**Series Standing Order ISBN 978-0-230-36085-3 (hardback)**

*(outside North America only)*

You can receive future titles in this series as they are published by placing a standing order. Please contact your bookseller or, in case of difficulty, write to us at the address below with your name and address, the title of the series and the ISBN quoted above.

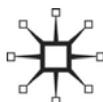
Customer Services Department, Macmillan Distribution Ltd, Houndmills, Basingstoke,  
Hampshire RG21 6XS, England

---

# Metaepistemology and Relativism

J. Adam Carter  
*University of Edinburgh*

palgrave  
macmillan



© J. Adam Carter 2016

Softcover reprint of the hardcover 1st edition 2016 978-1-137-33663-7

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

No portion 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, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

Any person who does any unauthorized 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 2016 by  
PALGRAVE MACMILLAN

Palgrave Macmillan in the UK is an imprint of Macmillan Publishers Limited, registered in England, company number 785998, of Houndmills, Basingstoke, Hampshire RG21 6XS.

Palgrave Macmillan in the US is a division of St Martin's Press LLC, 175 Fifth Avenue, New York, NY 10010.

Palgrave Macmillan is the global academic imprint of the above companies and has companies and representatives throughout the world.

Palgrave® and Macmillan® are registered trademarks in the United States, the United Kingdom, Europe and other countries.

ISBN 978-1-349-67375-9

ISBN 978-1-137-33664-4 (eBook)

DOI 10.1057/9781137336644

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources. Logging, pulping and manufacturing processes are expected to conform to the environmental regulations of the country of origin.

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

Library of Congress Cataloging-in-Publication Data

Carter, J. Adam, 1980–

Metaepistemology and relativism / J. Adam Carter, University of  
Edinburgh.

pages cm

Includes bibliographical references.

1. Relativity. 2. Knowledge, Theory of. I. Title.

BD221.C283 2016

121—dc23

2015032826

*For Emma*



# Contents

<i>Preface and Acknowledgements</i>	viii
1 Metaepistemology and Realism	1
2 Global Relativism	31
3 The Pyrrhonian Argument for Epistemic Relativism	59
4 Dialogic Arguments for Epistemic Relativism	77
5 Incommensurability, Circularity and Epistemic Relativism	107
6 Replacement Relativism: Boghossian, Kusch and Wright	137
7 A Different Kind of Epistemic Relativism	164
8 New Relativism: Epistemic Aftermath	192
9 Metaepistemology and Relativism	212
<i>Notes</i>	233
<i>References</i>	279
<i>Index</i>	293

# Preface and Acknowledgements

A key idea of Wittgenstein's in *On Certainty* that has always stood out to me as right is that in order to rationally investigate certain things, certain *other* things must already be in place in the background, and these other things in the background are simply taken for granted. The particular sort of thing I (*qua* epistemologist) usually end up investigating is what the epistemological facts are, and so maybe this seems a bit paradoxical. After all, the father of epistemology, Descartes, begins the *Meditations* by teaching us by example *not* to take things for granted – and this is a point that's hard to miss the deeper Descartes digs downward, questioning what we'd ordinarily take to be completely and utterly obvious.

But perhaps Descartes was already taking various general things for granted when doing *this* – when employing his method of doubt in the service of seeking sure foundations for human knowledge. And maybe epistemologists nowadays are taking broadly similar sorts of general things for granted, specifically, things about the kind of *answers we're looking for*, when following in his footsteps. Even more, it's not inconceivable that what it is epistemologists typically take for granted *about* the kinds of answers they're seeking when inquiring about such things as knowledge, justification, rationality, etc. is mistaken. In order to find out, we just have to take a step back from doing 'normal' epistemology in order to get epistemology's presuppositions in view – *viz.*, in order to think *about* the background itself.

The principal topic of this book is the relationship between contemporary (mainstream) epistemology and a certain view – relativism – which is a provocative kind of thesis *about* the very nature of the sort of answers – *viz.*, about *epistemological facts* – which mainstream epistemologists working today take one another to be attempting to uncover when writing, thinking and talking as they typically do, about the things they typically write, think and talk about.

Relativism is a very old and often reviled idea, but it never really died (though at various points in history many have tried to kill it off). In fact, as we'll see, the past decade in particular has seen relativist proposals take on a whole new life and shape. Yet, despite its long lineage, relativism is an idea that's difficult to get one's head around. To get one's feet wet with the idea of relativism, at least as it pertains to the kinds of

issues that epistemologists have a clear stake in, a run-of-the-mill epistemological dispute will be illustrative. Take, as a representative example, the following case which has been riling up epistemology experts and students alike since 1976:

**BARN FAÇADE:** Using his reliable perceptual faculties, Barney non-inferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact papier-mâché barn façades, erected as part of a plan to make the area he's in appear more prosperous (adapted from Pritchard 2012).

*Question:* Does Barney, in the above case, *know* that what he's looking at is a barn? What do you think? For readers not familiar with this particular case, it might already seem like a trick question: it seems we'd need to know a bit more about Barney. Was he drunk? Has his vision been tested recently? etc. While these are fair points, let's just assume that all sorts of things like this are covered and that everything here is just exactly as it would be in a normal case where we'd be happy to attribute knowledge to Barney when he looks directly at a barn. Suppose the *only* difference between a normal (good) case, and BARN FAÇADE, is that in BARN FAÇADE, there happen to be fake barns nearby, and moreover, that although Barney *didn't actually look at a fake barn* when forming his (true) belief that 'There is a barn', he very easily could have done so.

The question of interest to epistemologists (an interest that's remained more or less constant and at times heated for nearly 40 years now) is *whether* – and if so, *why* – the mere presence of nearby fakes should *matter* for whether Barney counts as knowing that he's looking at a barn. As things stand now, in 2015, the majority position in mainstream epistemology is that (propositional) knowledge is simply *not* compatible with the kind of epistemic luck that features in BARN FAÇADE. Though there remains a vocal minority of epistemologists who demur and take the opposite view. Interestingly, empirical results from experimental philosophy suggest that non-philosophers will be much more inclined to say, and contrary to what most epistemologists think, that Barney *does* know. So *who's actually right about this?* To put the rub of the question simply:

**Fake barn question:** *Is* (propositional) knowledge compatible with the kind of epistemic luck present in BARN FAÇADE, or is it *not*?

While 40 years haven't yet brought about unanimous agreement on the fake barn question (at least not yet), epistemologists on *both* sides behave just as if they are in agreement about a much more general point *about what they're disagreeing about* when reaching opposing verdicts on the fake barn question. In short, the shared background *agreement* on both sides seems to be that parties who disagree about how to analyse BARN FAÇADE simply can't *both* be right.

Enter here the *relativist*. The relativist looks at things here very differently. To a first approximation, the relativist is happy to say that 'yes' answers to the fake barn question can be true relative to some epistemic perspectives or standards, and 'no' answers to the fake barn question can be true relative to other epistemic perspectives or standards. But even more, the relativist wants to say that beyond these kinds of *relative truths*, there simply *is no further* perspective – or standard – independent sense in which 'yes' answers or 'no' answers to the fake barn question can aspire to correctness. There are *just* relative answers, nothing more.

And so (as the relativist sees it) in an *absolute* sense, you'll simply never – no matter how much epistemology you do – find out whether the chicken-sexer knows or whether virtue-epistemology can solve the Gettier problem because there are no absolute answers to these questions. The relativist relegates such questions to bad faith. This is, of course, a crude presentation (and oversimplification) of the relativist's view, but it will do for now.

If you feel a tension between what the relativist wants to say about epistemological facts and what mainstream epistemologists (e.g. the ones who keep arguing back and forth) seem to be *presupposing* about these very facts when they take themselves and their opponents to be trying to uncover them, then that's good. I think there is a real tension here. One thing epistemologists can do is to simply ignore the tension – or at least bracket it – and keep on answering questions like the fake barn question in the affirmative or negative, and then defending these answers as carefully as they can. (This is usually what I do, and I might add it's a lot of fun!)

Another thing to do is a bit more uncomfortable. The other thing is to take a step back and ask (in a crisis-of-conscience sort of way) whether mainstream epistemologist's knee-jerk tendency to proceed as though relativism can be brushed aside can ultimately be vindicated, or whether, instead, it turns out we should be taking the relativist more seriously than we actually do (at least, from within epistemology). The point of *Metaepistemology and Relativism* is to take this project on, and in doing so, to show where the complexities lie, to suggest why arguments

against more traditional kinds of epistemic relativism – once these arguments themselves are given their due – can't easily be redeployed against the new (semantic) variety, and finally (by the end) to suggest why, to the extent that epistemic relativism constitutes a threat to mainstream epistemology, it's probably a very different kind of threat than we might originally have been led to think.

The book is divided into nine chapters. Chapter 1, 'Metaepistemology and Realism', is a big picture chapter. The goal is to show what mainstream epistemologist's *metaepistemological* commitments are and where epistemic relativism stands in relation to these commitments. Since such commitments aren't often articulated by epistemologists (unlike in metaethics, where this kind of thing is explicitly talked about), I focus on the metaepistemological commitments that are revealed as the pragmatic presuppositions of paradigmatic first-order epistemological disputes. If the way I set things up in Chapter 1 is right, then an interesting result is that the question of the compatibility of mainstream epistemology's revealed metaepistemological commitments and epistemic relativism itself depends in a crucial respect on whether the arguments for epistemic relativism are any good in the first place – that is, whether the epistemic relativist's wider picture of epistemic facts is one we should embrace or reject. The task of evaluating the merits of argument strategies for epistemic relativism (in its traditional and more contemporary guises) consumes most of this book, from Chapters 2–8, before I connect the conclusions drawn on this score to the wider question about the relationship between metaepistemology and relativism in Chapter 9.

Chapter 2, 'Global Relativism', begins the investigation into the viability of epistemic relativism by considering whether there is any cause for taking a (very) quick road to epistemic relativism, one that proceeds through the comparatively more radical *global* version of the thesis. Global relativism, bizarre as it sounds, turns out to be much harder to dismiss than one would initially suspect, and thinking carefully about *why* this will be instructive in a number of ways which will be important later in the book (particularly, in Chapters 6 and 7) when engaging with the issue of how to best interpret what the epistemic relativist wants to say.

Chapters 3–5 engage with three specific and popular template argument strategies for motivating epistemic relativism. Chapter 3, 'The Pyrrhonian Argument for Epistemic Relativism', evaluates an argument strategy, pursued in recent work by Howard Sankey, which creatively redeploys the ancient Pyrrhonian problematic (traditionally a sceptical

argument) in the service of motivating epistemic relativism rather than scepticism. Chapter 4, 'Dialogic Arguments for Epistemic Relativism', takes on a general argument strategy which points to certain properties of actual (or possible) dialogues (e.g., the famous case of Galileo and Bellarmine) and concludes on the basis of the presence of these properties that epistemic relativism is true. Chapter 5, 'Incommensurability, Circularity and Epistemic Relativism', outlines and evaluates a familiar strategy-type for motivating epistemic relativism which draws from considerations to do with incommensurability and epistemic circularity. Spoiler: I don't think any of these strategy-types surveyed in Chapters 3–5 is compelling. A common theme I'll suggest is that that none of these argument strategies ultimately gives us a decisive reason to embrace relativism rather than scepticism.

That said, *even if* epistemic relativism could be philosophically motivated on the basis of the kinds of argument strategies surveyed in Chapters 3–5, there remains the separate but important issue of how to formulate the position in a satisfactory way. Chapter 6, 'Replacement Relativism: Boghossian, Kusch and Wright', engages with the plausibility of one popular semantic strategy for making sense of epistemic relativism – the *replacement model* – on which the relativist is interpreted as asking us to replace what must be rejected as absolutely false (e.g. unqualified claims of the form 'S's belief that p is justified') with, as Boghossian puts it, the 'nearest truths in the neighbourhood' which the relativist *can* accept, which are on the replacement model explicitly relational truths of the form 'According to epistemic system X, S's belief that p is justified.' Boghossian thinks the replacement model leads to incoherence. Kusch thinks, *contra* Boghossian, that there is a version of the replacement model that can be salvaged against Boghossian's criticisms, and Wright thinks that anyone who attributes the replacement model to the relativist has in doing so failed to take the relativist seriously, by taking seriously the idea that epistemic claims can be true or false 'albeit, relatively so'. On the interpretation of this standoff that I shall propose, we'll come to see that the core principle central to Burnyeat's (1976) reading of Plato's attempt to show Protagoras's global relativism to be self-refuting – what Burnyeat called the *principle of translation* (from Chapter 2) – will re-emerge as an insight about relative truth that is central to what is fundamentally at issue between Wright on the one hand, and Boghossian *and* Kusch on the other.

Chapter 7, 'A Different Kind of Epistemic Relativism', shows what epistemic relativism might look like if we leave the replacement model (and the associated principle of translation) behind and think about

relative truth in a very different way – where epistemic claims are ‘true, albeit, only relatively so’. This style of thinking about relativism has been championed in the main by John MacFarlane, whose brand of epistemic relativism takes the shape of an assessment-sensitive semantics for ‘knows’, according to which (for example) claims of the form ‘S knows that p’ get a truth value only relative to a context in which a use of the sentence ‘S knows that p’ is being assessed as true or false. After outlining MacFarlane’s rationale for giving ‘knows’ the relativist treatment he does, I raise several epistemologically grounded objections to MacFarlane’s project. In short, I show that MacFarlane-style relativism is in various ways at tension with a well-motivated position in epistemological metatheory called epistemic anti-individualism.

Chapter 8, ‘New Relativism: Epistemic Aftermath’, quickly qualifies the epistemologically oriented objections raised against MacFarlane’s version of epistemic relativism before developing a new dilemma for a proponent of a MacFarlane-style semantics for ‘knows’. The overarching move can be stated simply. I suggest, in a fashion that draws some close parallels with Allan Hazlett’s (2010) recent work on knowledge, factivity and knowledge ascriptions, that the *more* compelling MacFarlane’s argument is for his conclusion that the ordinary concept ‘knows’ is assessment-sensitive, the more reason the *epistemologist* has for thinking that the ordinary concept of knowledge is *epistemologically* uninteresting, and moreover, so are ordinary knowledge ascriptions. Crucial to the line advanced here will be an examination of what contemporary epistemology might look like if it were centred around an assessment-sensitive concept of knowledge.

Chapter 9, ‘Metaepistemology and Relativism’, has two overarching aims. First, the conclusions drawn from Chapters 2–8 will be situated within a wider context: that of the complicated relationship between metaepistemology, realism and relativism, with an eye to answering questions left outstanding at the end of Chapter 1. The second thing I want to do is to motivate, in light of what’s been argued for already, an entirely different way of thinking about the relationship between mainstream epistemology and the kind of challenge to it that arguments for epistemic relativism stand to pose. If I am right, at least some arguments for epistemic relativism (particularly, new epistemic relativism) *can*, and despite what I insinuate in Chapter 8, have an important kind of relevance to mainstream epistemology, even if more traditional arguments of the sort canvassed in Chapters 3–5 do not. Though the *kind* of relevance is – I’ll argue – of an entirely different sort than is ordinarily thought.

Thanks to Matthew Chrisman, James Collin, Bolesław Czarnecki, Pascal Engel, Mikkel Gerken, Steven Hales, Allan Hazlett, Jesper Kallestrup, Martin Kusch, Michela Massimi, Robin McKenna, Paul Noordhoof, Orestis Palermos, Christian Piller, Brian Rabern, Conny Rohde, Shane Ryan, Stephen Ryan, Markus Seidel, Claudine Tiercelin and Kevin Wallbridge and to audiences at the University of Edinburgh, University of Vienna, York University for helpful feedback, discussion and comments on draft chapters. Thanks also to Natalie Ashton for organising and leading a stimulating relativism reading group at Edinburgh in 2014 and to Martin Kusch and Robin McKenna for organising a series of talks on my book at the University of Vienna as part of Kusch's European Research Council funded 'Emergence of Relativism' grant. Special thanks go to Maria Baghramian for originally inspiring my interest in relativism and whose work has been important in shaping my thinking, to Duncan Pritchard for continual and unwavering support and encouragement throughout the planning and writing of the book, to Aidan McGlynn for detailed and helpful feedback on much of the book and to Modesto Gómez Alonso, whose enthusiasm and discussion, comments and conversation throughout the writing of the book was invaluable.

Extra special thanks to my amazingly supportive partner Emma Gordon for reading drafts of every chapter just as soon as I finished them, for talking through the main ideas with me, and for her loving and patient support (and great music playlists) during those hundreds of nights of drinking instant coffee and staying up too late.

J.A.C  
Edinburgh

# 1

## Metaepistemology and Realism

*Abstract.* Metaepistemological commitments are revealed in first-order practice, though (unlike in metaethics) are not often given explicit expression. This chapter does two central things. The first is to develop a reasoned way of locating revealed second-order commitments in metaepistemology by looking straight to paradigmatic first-order disagreements, and to what is common ground to these disagreements. The second is to show that a pervasive element of the common ground in first-order epistemological debates turns out to be a commitment to at least a minimal form of *metaepistemological realism*. I conclude by considering the kinds of claims that mainstream metaepistemology tacitly excludes, in virtue of presupposing this kind of realism, and how we might best locate epistemic relativism within the picture proposed.

### 1.1 Introduction

Metaethics, as Sayre-McCord (2014) aptly puts it, attempts to ‘step back from particular substantive debates within morality to ask about the views, assumptions, and commitments that are shared by those who engage in the debate’. Think of metaepistemology as doing the same thing, but for epistemology.<sup>1</sup> One striking difference between metaethics and metaepistemology is that there are a lot of people actually *doing* the former. Accordingly, metaethical commitments are often articulated, explicitly and carefully. With a few outlying exceptions, the same isn’t so for metaepistemology. Epistemologists by and large carry on in their first-order projects without taking the time to articulate the more general commitments operating in the background.

An examination of metaepistemology ought to distinguish between two sources of metaepistemological commitment: on the one hand, one can take on a metaepistemological commitment by simply articulating such a view (as the metaethicists often do). Call these *articulated* metaepistemological commitments. A separate source of metaepistemological commitment is unarticulated, but revealed in first-order practice; call metaepistemological commitments of this latter variety *revealed* metaepistemological commitments. Given the inchoate state of metaepistemology, the focus here will be primarily on the latter: that is, we'll look to the action in *first-order* debates in mainstream epistemology in the service of characterising what the *second-order* commitments are. The wider objective will be to argue for and characterise mainstream metaepistemology as a function of what characteristic debates in mainstream epistemology take for granted.

On first blush, we can envision two *prima facie* plausible ways to do this – *viz.*, to locate revealed metaepistemological commitments by looking straight to first-order debates, the important disputes actually going on. The first way proceeds as follows: we look to some of the most striking disputes in first-order epistemology with an eye to mapping these first-order disputes on to more general second-order dividing lines. This is more or less the strategy that has been pursued by William Alston (1978) and Jonathan Dancy (1982) in their early attempts to carve up metaepistemological space.

An alternative strategy starts in the same place – *viz.*, first-order debates in epistemology – but, rather than trying to map first-order disagreements on to second-order disagreements, the second strategy we can envision tries to reason from these first-order disagreements to background *second-order agreements* – that is, by locating the common ground lying behind first-order disagreements and with reference to which we can best explain why first-order disagreements take the characteristic *shapes* that they do in practice. I'll be opting for the second strategy in what follows, and in doing so, it will be shown further how what's taken as common ground connects with certain features of realism; the more general picture I'll develop here is one where participants in mainstream epistemology's paradigmatic debates can be understood as by and large revealing a tacit background commitment to at least a minimal form of metaepistemological realism.

Here is the plan. In §1.2, I'll show why Alston's and Dancy's respective attempts to map first-order disagreements on to *second-order* disagreements in epistemology aren't compelling. In §1.3, I'll sketch an argument for doing things the other way around. In particular, I'll outline a plausible picture on which the notion of presupposition – and in particular

*pragmatic presupposition* – can be connected to both metaepistemological commitment and as well as to disagreement; what emerges is a rationale for locating metaepistemological commitments in the *common ground* of first-order disputes. In §1.4 I sharpen this picture by using, as a case study, the perennial debate between Moore and the sceptic; with reference to the Stalnaker/Grice model of common ground, we'll see that Moore and the sceptic are (despite little *first-order* agreement) presupposing a certain shared background commitment to *epistemic facts* with an objective profile. And from Moore and the sceptic we can generalise. The upshot of §§1.2–1.4 will be that paradigmatic disputes in mainstream epistemology (much like the one between Moore and the sceptic) betray a revealed metaepistemological commitment to – by specifically taking for granted – the objectivity of epistemic facts under discussion. The remainder of the chapter shows how the particular picture of epistemic facts we find in the common ground (of paradigmatic first-order debates) *de facto* satisfies plausible conditions of a *realist* picture of epistemic facts. Toward this end, the approach taken in §1.5–1.6 defines and clarifies a kind of *generic realism* – as a combination of (suitably articulated) *existence* and *independence* theses, and defines (with reference to this working general picture of realism) *metaepistemological realism*, as constituted by a conjunction of existence and independence theses about epistemic facts. In §1.7 I show how mainstream metaepistemology's commitment to epistemic facts (as developed and defended in §§1.2–1.4) constitutes a minimal form of metaepistemological realism – *mainstream metaepistemological realism*. In §1.8 I conclude by charting a range of metaepistemological *anti-realist* views that are precluded by the kind of realist picture first-order epistemologists take for granted.

## 1.2 Metaepistemological dividing lines: two early attempts

### 1.2.1 Alston and the 'fact/value' divide

Let's look first at how Alston tried to carve out what he took to be the most interesting metaepistemological dividing lines. Alston, it is crucial to note, understood the *metaethical* landscape (at least, in the late 1970s), as featuring three prominent positions.

- (i) *metaethical noncognitivism*, according to which ethical judgements – (e.g. Murder is wrong) – do not express truth-apt beliefs;<sup>2</sup> and, among cognitivist positions,

where such judgements are regarded as truth-apt:

- (ii) *metaethical naturalism*, according to which – at least, as Alston characterises it – ‘ethical terms (concepts, statements) can be defined, explicated, or analyzed in “factual” terms, that they, at bottom are factual terms, and that ethical questions are, at bottom, questions of fact, perhaps of an especially complicated sort’ (Alston 1978, 276).
- (iii) *metaethical intuitionism*, which ‘maintains that ethical concepts are *sui generis*, that they are of a distinctively and irreducibly normative or evaluative sort, not to be reduced to matters of fact, however complex’<sup>3</sup> (*ibid.*, 276).

With reference to these three positions, Alston drew three corresponding metaepistemological conclusions:

- (C1) that parallels with non-cognitivism in epistemology are rare;
- (C2) that the epistemological parallel to ethical naturalism features in the background of the ‘causal theory of knowledge’<sup>4</sup>, according to which (roughly), *S* knows that *p* just in case *S*’s belief that *p* has a certain kind of causal history; and that
- (C3) the epistemological parallel to ethical intuitionism features in the background of the JTB-style analysis of knowledge,<sup>5</sup> where ‘justified’ is taken as an evaluative term.

Regarding the first of these three conclusions: Alston simply didn’t anticipate, at the time of writing, anything like contemporary versions of epistemic expressivism, which has been explored only in recent decades,<sup>6</sup> but we can set this point aside. C2 and C3 reveal how Alston envisions the key metaepistemological dividing lines: there is *metaepistemological naturalism* (mapping on to the causal theory) and *metaepistemological intuitionism* (mapping on to the JTB theory).

Obviously, a *prima facie* objection already to this picture is that Alston – by focusing on just two kinds of rather specific positions in epistemology – was oversimplifying the contemporary epistemological landscape. This charge is fair, and we shall return to it. But for now, this oversimplification will be useful. It provides us a simple way to think about metaepistemological commitments as (at least, potentially) in connection with first-order commitments.

That said, it won’t be hard to see how the simple picture Alston gives us faces some problems. Consider (C2) and (C3). By these conclusions, Alston is committed to saying that (for example) Alvin Goldman (an arch-promoter of the most prominent contemporary version of the

causal theory, *reliabilism*<sup>7</sup>) and Richard Feldman (an arch-promoter of a non-causal, internalist variety of the JTB account, *evidentialism*), are espousing not only different positions in (first-order) epistemology, but moreover, that they are importantly divided along *second-order* lines. Specifically, Alston must grant that (*qua* proponent of a causal theory) Goldman would be in a position to claim a remarkable advantage over Feldman's evidentialist-variation of the JTB theory<sup>8</sup> – namely, that his causal theory offers a theory about something, justification, which 'can be defined, explicated, or analysed in 'factual' terms' (1978, 276) whereas Feldman, in virtue of endorsing a variant of the JTB-view, is not.<sup>9</sup> Unsurprisingly, this is not the way causal theorists of justification and knowledge actually in practice attempt to distinguish themselves from their opponents.

Furthermore, if Alston's assessments in (C2) and (C3) were correct, then Feldman would find himself in a position to simply shift the focus to a different, (perhaps equally) remarkable, advantage that non-causal JTB accounts would have in comparison to causal theories. Specifically, Feldman could, by Alston's lights, rightly insist that his favoured position – *viz.*, *evidentialism* – is distinguished as a genuine theory of epistemic justification (or, justifiedness<sup>10</sup>), but – despite what Goldman professes – Goldman's reliabilism, in virtue of its causal core, is not.<sup>11</sup> This is clearly wrong.

Goldman and Feldman are at odds about a lot in epistemology. But they at least take each other to be trying to do the same *kind of thing* when theorising about epistemic justification and knowledge. That's one of the reasons they have debated back and forth, and why their debates make sense. Compare: their behaviour over the past 30 years would be utterly mysterious if they did *not* take each other to be doing (as regards any alleged fact/value divide which Alston (1978, 276) is adverting to in distinguishing naturalism and intuitionism) the same *kind of thing*. If indeed there is an interesting fact/value dividing line in metaepistemology, causal theorists and non-causal JTB theorists at least behave as though they regard each other to be on the same side of it.<sup>12</sup>

### 1.2.2 Dancy and the 'monism/pluralism' divide

Jonathan Dancy, writing in the early 1980s, thinks – like Alston did – that the first-order divide between causal theories and non-causal JTB theories (of knowledge and epistemic justification) directs us to an interesting metaepistemological division. However, as we'll see, Dancy struggles to make the case as well, though for different reasons. Dancy (1982), like Alston, takes his objective to be that of drawing fruitful connections

between metaethics and metaepistemology, and he takes as a starting point that, in metaethics, ‘the most revealing criterion in the dispute between naturalist and intuitionist’<sup>13</sup> is not (as Alston had thought) one between what’s reducible to matters of fact and what’s not; rather, it is between *monism* and *pluralism*.<sup>14</sup>

Dancy characterises two features of the ‘monism’ exhibited by metaethical naturalists: firstly, *metaethical naturalists*, in virtue of their *monism*, find a single property in virtue of which all right actions are right; and secondly, they are committed to the view that ethical principles can be derived from a common principle. *Intuitionists*, by their commitment to *pluralism*, insist firstly that there are at the least several different properties which make right actions right; and, secondly, that there are multiple ethical principles which cannot be derived from a common principle.<sup>15</sup>

As Dancy envisions the metaethics/metaepistemology parallel, causal theories in epistemology – in a manner congruous with metaethical naturalists – attempt to find a ‘common pattern or shape in all justified beliefs’, whereas JTB-theorists who reject the causal theory may (in a way relevantly analogous to ethical pluralism) ‘do so on the grounds that...beliefs can be justified in other ways as well’ (Dancy (1982, 399)). This characterisation, though, turns out to be misleading in several respects.

Are ‘causal theorists’ (e.g. reliabilists) really ‘monists’ in a way that constitutes an interesting departure from non-causal JTB theorists? Dancy thinks so because he thinks that, on causal theories, the ‘natural bases’ from which justified beliefs are justified bear some single property. This claim is suspect, but let’s bracket it.<sup>16</sup> A more serious worry with Dancy’s diagnosis is revealed by his insistence that non-causal JTB theorists should be thought of as ‘pluralists’ (in the sense Dancy intends) in a way that causal theorists are not. Dancy (1982, 398) writes that:

Chisholm, a leading JTB-theorist, looks like a pluralist. He offers a large number of different epistemic principles (principles of justification) without attempting to derive them from a common source.<sup>17</sup>

As a preliminary point, it is not at all clear why causal theorists like Goldman – or, for that matter, any epistemologist who adverts to the notion of causation in attempting to illuminate knowledge<sup>18</sup> – should be understood as attempting to derive all principles relevant to epistemic justification from a ‘common source’. (Though Mill did so in ethics,

there's really no analogue in contemporary epistemology.) Consider, after all, that the bare statement of process reliabilism is not derivable from (for instance) Goldman's distinct but complementary *principle of veritistic value*, according to which the epistemic value of processes and practices is to be assessed on the basis of its promotion of true beliefs and avoidance of error.<sup>19</sup> But neither is the principle of veritistic value orthogonal to Goldman's reliabilist account of justification; the principles are complementary but not interderivable.

Even more problematically, though, it's not so clear why (non-causal) JTB theories should be regarded as 'shapeless'<sup>20</sup> in a way that that causal theories are not. JTB projects share a commitment to searching for 'conditions that are individually necessary and jointly sufficient for propositional knowledge'.<sup>21</sup> But this is also exactly what causal theories are, in the main, and self-confessedly, attempting to do.<sup>22</sup> Any *interesting* metaepistemological dividing line between causal and non-causal theories is thus not plausibly going to be drawn between 'monism' and 'pluralism' any more than it is between 'fact' and 'value.'

In short: examining Alston and Dancy confirms what may have already seemed very plausible, namely *that first order disagreements aren't good indicators of second-order disagreements*. In what follows, I'll suggest we take this idea a step further and say that things are more or less the other way around: established first-order disagreements (like the disagreement between externalists and internalists) are reliable indicators of certain kinds of *second-order agreements*.<sup>23</sup> I want to now pursue this idea in some more detail – which will involve putting several pieces together – and then to show how, eventually, it connects to the issue of what epistemological realism might look like.

### 1.3 Metaepistemology, presupposition and disagreement

One reason first-order disputes didn't map on to second-order disputes very well is that in order to get first-order disputes up and running, a range of things must already be in place in the background, by both sides. This, as we saw, certainly seemed to be the case with Goldman and Feldman. Let's now take this idea a bit further. In this section, I'll sketch a picture on which established first-order disputes hold a special key to locating second-order (metaepistemological) agreement. Once the picture is in place, I'll then put it to work in the service of assessing (by looking at mainstream debates) just *what* some of the most fundamental metaepistemological commitments are. A central element of the simple picture I'll put forward will be the notion of a *presupposition* – which

will function as the ‘bridge’ between metaepistemological commitments and first-order disagreements.

The first element of this simple view will involve ‘linking together’ metaepistemological commitment and presupposition. We can do this relatively straightforwardly, in a way that is more or less intimated by a Sayre-McCord-style gloss of metaepistemology,<sup>24</sup> as the *presuppositions (and commitments of these presuppositions) of first-order epistemology*.

**Presupposition/Metaepistemological Commitment (PMEC):** One’s *metaepistemological commitments* will be (at least, in part<sup>25</sup>) a matter of what one’s first-order projects *presuppose*, and the commitments of these presuppositions.

PMEC is a simple and attractive way to think about revealed metaepistemological commitment.<sup>26</sup> PMEC says, albeit at some level of generality, *where* metaepistemological commitments are revealed: they are revealed in what first-order projects presuppose. So: *how do we work out what first-order disagreements presuppose?*

Here we need a second element of the view – one that provides some principled way of connecting the kind of presuppositions that constitute metaepistemological commitments with first-order disagreements – so that we can look straight to the latter and actually locate the former.

This leg of the view – which will involve connecting presupposition with disagreement – will take a bit of work and then some refinement. Firstly, regarding disagreement. It’s beyond what we can do here to defend necessary and sufficient conditions for disagreements – (e.g. like the one between Feldman and Goldman about justification and knowledge) – to be *genuine*, as opposed to, say ‘merely verbal’.<sup>27</sup> But for our purposes, it will suffice to highlight two simple characteristics of paradigmatic first-order disagreements (e.g. when for some epistemological proposition *p*, *S*<sub>1</sub> affirms *p* and *S*<sub>2</sub> denies it) that are relatively uncontroversial. The first is *semantic*, the second *pragmatic*. The *semantic* characteristic is that *S*<sub>1</sub> disagrees with *S*<sub>2</sub>’s belief that *p* only if *S*<sub>1</sub> has beliefs with contents incompatible with *p*.<sup>28</sup> The *pragmatic* characteristic is that genuine disagreements typically feature certain patterns of linguistic data (e.g. explicit acknowledgment of contradiction, etc.).

Correspondingly, (and rather conveniently) there are two central notions of presupposition, *semantic presupposition* (e.g. in the Frege–Strawson tradition) and *pragmatic presupposition* (in the Grice–Stalnaker

tradition). Naturally, then, there are two ways that we might be inclined to connect disagreement to presupposition. One way will be to connect the ‘semantic conflict’ feature of disagreement with the notion of semantic presupposition; another will be to connect the pragmatic feature of disagreement with the notion of pragmatic presupposition.<sup>29</sup> For our purposes of understanding metaepistemological commitments by investigating first-order disagreements, the latter connection is going to be more appropriate and interesting than the former will be – or so I’ll soon suggest.

But first, a few quick remarks about the former connection – between semantic conflict and semantic presupposition. Semantic presupposition (in the Frege–Strawson tradition) concerns what sentences presuppose, and on this view, one sentence presupposes another if, and only if, whenever the first is true *or* false, the second is true.<sup>30</sup> Now, if we advert to this semantic constraint on disagreement according to which disagreement entails incompatible contents, and combine this with something like a simple semantic account of presupposition, according to which a shared presupposition is going to be entailed by the assertion of both a content and its denial, then it follows that any time two individuals disagree *vis-à-vis*  $p$ , then there will be at least something  $\sigma$  presupposed by (at least, the sentences used by) both parties to the disagreement. And here we can envision how at least some space would open up for trying to locate metaepistemological commitments by looking straight to first-order disagreements and, specifically, to the semantic presuppositions of the sentences characteristic of these disagreements.

This semantic strategy – on which the semantic conflict feature of disagreement is connected to the notion of semantic presupposition – would not only be tedious, but it would seem (for several reasons) to put the focus in the wrong place;<sup>31</sup> the most obvious strike against trying to elucidate metaepistemological commitment *via* the semantic presuppositions of first-order disagreements is that the kind of presupposition that seems most relevant to characterising metaepistemological commitment does not have entailment at its heart, but something else.

Fortunately, there is a much more relevant notion of presupposition – *pragmatic presupposition* – that we can easily connect to both PMEC and to features of first-order disagreements. If we think of the proper object of the semantic notion of presupposition as what words or sentences presuppose, we can think of the proper object of *pragmatic presupposition* as what people take for granted when speaking.<sup>32</sup> The notion of pragmatic presupposition has been pioneered most influentially by Robert

Stalnaker (1973), who himself drew heavily from Grice. As Stalnaker (2002, 701) sums up the crux of the Stalnaker/Grice view:

To [pragmatically] presuppose something is to take it for granted, or at least *to act as if one takes it for granted*, as background information – as common ground among the participants in the conversation.<sup>33</sup>

We'll get to common ground shortly. But first, let's register now the second leg of our simple model – the leg connecting presupposition with first-order disagreement.

**Pragmatic Presupposition (Stalnaker/Grice):** Pragmatic presuppositions of first-order disagreements are the common ground in such disagreements, as reflected through behaviour and use of language by participants to the disagreements.

We can now simply *combine* PMEC with the Stalnaker/Grice account of pragmatic presupposition to get the result that: *If disagreements at the first order in epistemology are ones where both sides to the disputes are disposed to behave, in their use of language, as if they believe some metaepistemic claim  $\sigma$  to be common ground in the context of their dispute, then  $\sigma$  is a metaepistemological commitment of both sides of the first-order dispute.*

This result, from combining PMEC with the Stalnaker/Grice account of presupposition, offers a useful working picture that tells us how to find metaepistemological commitments by looking to first-order disagreements. Let's now turn to putting this simple view to work.

#### 1.4 Common ground: a lesson from Moore and the sceptic

Four things about common ground. Firstly, what is common ground between *A* and *B* is ordinarily *believed* by *A* and *B*.<sup>34</sup> Secondly, *p* is common ground in disagreement between *A* and *B* only if a certain iterative condition is in place, according to which each takes for granted that the other believes that *p*, and so on. Thirdly, what is common ground between *A* and *B* can be false. Fourthly – and this will take a bit more unpacking – what counts as common ground is, on the Stalnaker/Grice model of presupposition, 'reflected in practice', by how the participants to the dispute behave, through their use of language.

To illustrate this latter idea, it will be helpful to consider – as an instructive example – a first-order debate so dialectically entrenched on both sides that there seems to be little common ground at all. Here I

have in mind the classic debate in 20th century epistemology between an epistemological sceptic and a Moorean.<sup>35</sup> The rationale for looking here is that finding common ground in the especially wide gap between Moore and the sceptic will make it easier to find it elsewhere, and more generally.

In short, the neo-Moorean responds to the radical sceptical challenge (cast in terms of a denial that we can know that there is an external world) by *a la* Moore, *turning the tables*.<sup>36</sup> That is, the Moorean asserts a few commonsense things she knows – the first premise of her argument – and then reasons (e.g. *via* closure<sup>37</sup>) that if she knows these things, then she knows *ipso facto* that there is an external world,<sup>38</sup> which of course the sceptic insists Moore does not know – nor indeed anyone else for that matter. In this deep and enduring dispute, each side is – when engaged with the other – often accused of begging the question.<sup>39</sup> Framed in terms of a simple disagreement: let  $K(W)$  represent the proposition that *Moore knows there is an external world* – a proposition the Moorean affirms and her opponent denies.<sup>40</sup> *Vis-à-vis* this dispute, what is the *common ground* reflected in practice?

It's tempting to say 'nothing!' After all, as Keith Lehrer (1971) and Barry Stroud (1984) have observed, it looks illicit for Moore ever to have begun as he did in this context, by claiming to have known his premise in the first place.<sup>41</sup> A more general worry is that whether Moore knows something – *viz.*, that there is an external world – something that his own premise entails, is precisely what's at issue.

But it would be too quick to say there is simply *no* common ground here. Moore and the sceptic are clashing about the *scope* of human knowledge, but to do so, they are taking for granted, for starters, certain things about its *nature*. Consider: both Moore and the sceptic are taking for granted that there are certain constitutive principles governing what counts as correct application of the concept of 'knowledge'. One such principle is that 'knowledge' requires truth (e.g.  $S$  knows that  $p$  only if  $p$  is true). Without this in the shared background, for instance, we'd expect the disagreement to reflect what would be a *much weaker position for the sceptic*. In particular, we'd have a hard time making sense of the sceptic's challenging Moore's claim to know his first premise.<sup>42</sup> Moore and the sceptic are thus revealing, in practice, that a truth condition on knowledge is shared common ground between them. Their debate would take a *very different shape* without it.

I think we can helpfully sharpen this rationale, in terms of the notion of *sensitivity*. We've just reasoned counterfactually that: *if the denial of the claim in question* (that knowledge is factive) *were assumed by both parties*

holding everything else fixed – i.e. that is, in the closest situation where the operative assumption is that knowledge did *not* entail truth – then the debate between Moore and the sceptic would not simply continue on as it does in the actual world, but rather, it would take on a very different shape, in practice – *viz.*, the sceptic would not be in a position to, like Lehrer and Stroud do, object to Moore claiming to know his first premise. And this much was taken as an intuitive explanation for why the factivity of knowledge is plausibly taken as part of their common ground. In this respect, it seems like the following kind of ‘sensitivity principle’ looks very plausible as a working guide as to what is being assumed in the common ground.

**Sensitivity Principle for Common Ground (SPCG):** For  $A$ ,  $B$  and dispute  $D$ ,  $\sigma$  is part of the common ground between  $A$  and  $B$  *vis-à-vis*  $D$  if the following counterfactual is true: were not- $\sigma$  assumed by  $A$  and  $B$ ,  $D$  would be relevantly<sup>43</sup> different.

While ‘relevantly’ here is of course a weasel word, something like SPCG is surely right, insofar as we want to take seriously Stalnaker’s suggestion that common ground is *reflected* in practice; after all, if practice were not sensitive to items in the common ground, as claimed by SPCG, then common ground wouldn’t be reflected in practice (rather: practice would be insensitive to what is in the common ground).

Returning to Moore and the sceptic: with reference to SPCG, there’s a straightforward way to capture the initial point that the factivity of knowledge is common ground in their debate. But that said, this *particular* item of common ground is an example of common ground that is specific to disputes about knowledge. (Debates about the structure of intellectual virtues, for instance, needn’t include the factivity of knowledge in the common ground.) Naturally, the more interesting items of common ground, at least for the purpose of identifying more general metaepistemological dividing lines, will be ones that are shared in equal measure *across* other paradigmatic debates in epistemology, including debates that aren’t about knowledge at all.

To this end, let’s consider, at a greater level of generality, (and with SPCG in hand) two more elements of the common ground between the Moorean and the sceptic. The first involves epistemic facthood, and the second involves something more specific about the character of epistemic facts.

Now Moore, just as well as his opponent, must already assume – in engaging in debate as they do – that there is some fact of the matter – *viz.*,

that Moore either does, or does not, instantiate a certain epistemic property – in this case, the property of knowing that there is an external world. To get a feel for why a tacit commitment on both sides to (regarding each other as embracing) epistemic facthood lies in the background of their dispute, just consider how this commitment is actually betrayed in practice.

With reference to SPCG, let's simply plug in the dispute between Moore and the sceptic about whether  $K(W)$  for 'D' and further let  $\sigma$  be the claim that *there's a fact of the matter vis-à-vis whether  $K(W)$* . SPCG tells us that the claim that there's a fact of the matter whether  $K(W)$  is in the common ground between Moore and the sceptic, *vis-à-vis* the dispute about whether  $K(W)$ , provided the following counterfactual is true: were both parties to assume there was no fact of the matter whether  $K(W)$ , their dispute would be relevantly different, as revealed in practice.

Now, let's simply imagine this. How *would* Moore and the sceptic proceed, *vis-à-vis* whether  $K(W)$ , in the nearest worlds where (holding everything else fixed) they assume (and assume each other assumes, etc.) there is *no fact of the matter* about whether  $K(W)$ . It's tempting to say that they simply wouldn't continue to assert their conclusions. Though I think we can be a bit more charitable. It might well be that in such worlds, Moore and the sceptic continue to *say* their conclusions,  $K(W)$  and  $\neg K(W)$ , respectively. After all, perhaps the nearest worlds where they take it that there is no epistemic fact of the matter whether  $K(W)$  are nonetheless worlds where they continue to express some pro- or con-attitudes, and use surface-grammar declarative sentences to do this. *But*, even if we grant that Moore and the sceptic will continue to proclaim  $K(W)$  and  $\neg K(W)$  in such worlds, Moore will *not*, in such a world, offer as he does a *proof* for  $K(W)$ <sup>44</sup> – where a proof is something Moore takes as industry standard in the service of *settling questions*. Here's Moore:

...I do want to emphasize that, so far as I can see, we all of us do constantly take proofs of [the sort Moore offers the sceptic] as absolutely conclusive proofs of certain conclusions – *as finally settling questions*, as to which we were previously in doubt.<sup>45</sup> (Moore and Lewy (1962, 167), my emphasis)<sup>46</sup>

In the nearest worlds where neither Moore nor his sceptical opponent are taking for granted that there are epistemic facts, but rather, taking for granted that there are *not*, Moore simply does not aim with any plausible communicative purpose to *establish* his conclusion, beyond doubt, in the manner that he does.<sup>47</sup> To make sense of their argument,

we must take them to be assuming (and assuming each other assumes) that there are epistemic facts – in this case, an epistemic fact consisting *either* in Moore instantiating (as Moore thinks), or not instantiating (as the sceptic thinks), a certain epistemic property.<sup>48</sup>

Now that said, I'd like to suggest further that Moore and the sceptic are also jointly taking for granted something a bit deeper about the *character* of epistemic facts (such as the salient one they're disputing). Put simply, they are both taking it for granted that whatever the epistemic facts are, they hold *equally for Moore just as well as for the sceptic* or anyone else and, as such, are not *merely* facts that hold just for one, or for the other. With reference to SPCG, *were* Moore and the sceptic assuming otherwise, they would not take their respective positions to *preclude* one another. But they do take their respective positions to preclude one another. As Keith DeRose (2004) remarks:

...the sceptic and her opponent...take themselves to be contradicting one another; each intends to be contradicting what the other is saying; and, beyond what's going on privately in their own minds, each is publicly indicating that they are (or at least mean to be) contradicting the other, by saying such things as, 'No, you're wrong. I do know.'<sup>49</sup> (DeRose (2004, 3))

It is against a background commitment to thinking that epistemic facts (such as the one under dispute) are universal (hold for both alike) and they hold independently of Moore's or the sceptic's own perspectives – call this the *objectivity* of epistemic facts – that this mutual recognition of contradiction makes sense in practice, and why Moore sincerely claimed to think the sceptic's position could be 'shown to be wrong'<sup>50</sup> (Moore (1939, 25:148)).

Let's abstract now from the situation between Moore and the sceptic. If Moore and the sceptic, in a dispute whereby they are accused of begging the question against one another, nonetheless have as common ground a commitment to objective epistemic facts (in the sense described above), then we can see how, very plausibly, the same goes for Goldman and Feldman. And as well for coherentists and foundationalists about the structure of justification, for those for and against factive reasons, for those for and against pragmatic encroachment, and so on.

In sum: the disagreements we find *throughout* first-order epistemology pragmatically presuppose (by interlocutors on both sides of the key issues) a commitment to taking for granted epistemic facts which hold equally for both disagreeing interlocutors – *viz.*, epistemic facts with an

objective profile – by revealing this much in practice to be part of the common ground that must already be in place for their behaviour to make sense. Thus, what results is something like:

- (1) Most first-order disagreements in epistemology are ones where both sides to the disputes are disposed to behave, in their use of language, as if they believe a commitment to epistemic facts with an objective profile is common ground in the context of their dispute

Let's now combine this result to the result that emerged from combining PMEC with the Stalnaker/Grice account of pragmatic presupposition:

- (2) *If disagreements at the first order in epistemology are ones where both sides to the disputes are disposed to behave, in their use of language, as if they believe some metaepistemic claim  $\sigma$  to be common ground in the context of their dispute, then  $\sigma$  is a metaepistemological commitment of both sides of the first-order dispute.* (From MPEC, and Stalnaker/Grice pragmatic presupposition).

What results is:

- (3) **Mainstream metaepistemology (MM):** A metaepistemological commitment of most first order-disagreements in epistemology is: a commitment to taking for granted epistemic facts with an objective profile.

## 1.5 Realism: generic and metaepistemological

A quick interlude, before we connect MM to the issue of metaepistemological *realism*. Note that I have not suggested (at least not yet) that there are any epistemic facts; nor have I suggested that if there were epistemic facts, that they would be objective in character. I've not even gone so far as to suggest that the kinds of disputes we find at the first order, in virtue of pragmatically presupposing objective epistemic facts in the common ground, provide us with evidence from which we can reasonably infer that there are objective epistemic facts. It might after all be that most first-order disagreements in epistemology pragmatically presuppose something that is categorically false, perhaps even *way* off the mark.<sup>51</sup>

Accordingly, and to emphasise, what has been suggested to this point is perfectly compatible with the claim that (as, for instance, the

metaethical noncognitivist would tell us) apparent disagreements arise even when there are no facts to disagree about. After all, the noncognitivist does not deny that there is a position that is metaethical realism, a position characterised in part by a commitment to the kind of objective moral facts that many normative ethicists think they're arguing about. (Allan Gibbard, for instance, grants this much.<sup>52</sup>) What the noncognitivist thinks is that metaethical realism, *in light of its commitments to such facts*,<sup>53</sup> should be rejected.

But, if we told the metaethical noncognitivist that, in epistemology, while folks don't typically articulate explicitly anything like 'metaepistemological realism',<sup>54</sup> it's nonetheless the case that epistemic facts with an objective profile are part of the common ground that is needed to make sense of how most all first-order debates actually proceed, what would she say? I think she would likely tell us that what we've got in epistemology is a bunch of people committed *tacitly* to (at least some version of) realism. And I think she'd be right. I want to now clarify how mainstream epistemology's revealed metaepistemological commitments involve a wider commitment to a kind of metaepistemological *realism*.

Toward this end, let's put on the table a simple working idea of what *generic realism* – that is, realism in any area of discourse – plausibly amounts to. In doing so, we'll see more clearly how, in any given area  $\alpha$ , a commitment to  $\alpha$ -facts is relevant to whether or not one is committed to a form  $\alpha$ -realism. Looking at realism in this abstract way will put us in a good position to see what has to be in place for a view to count as a version of *metaepistemological realism*. In doing so, we'll see exactly how the endorsement of epistemic facts with an objective profile (that is presupposed by disputes in mainstream epistemology) tacitly commits (most) mainstream epistemological projects to a kind of metaepistemological realism.

### 1.5.1 Generic realism

It is notoriously difficult to say what, for some subject matter  $\alpha$ , realism *vis-à-vis*  $\alpha$  is. One guiding thought, noted by Simon Blackburn (1993), is that an  $\alpha$ -realist takes  $\alpha$ -claims to be validated by the way things stand in the world. Most of us are realists, for example, when it comes to cat discourse; we think that claims about the Tonkinese cat in the next room will be validated by how things stand. Validated in that, we think, there really are states of affairs consisting of cats existing and bearing the properties we ordinarily attribute to them – rendering some of our statements about them true. To think our statements about cats are validated by how things stand intimates how it is that we think the correctness of

cat discourse answers to some way the world (with cats and their properties) really is – as opposed to, say, how we merely hope, imagine, agree or perceive that it is.

What falls naturally out of Blackburn's simple 'validated by how things stand' locution is two distinct – i.e. logically independent – kinds of commitment. One kind of commitment involves *existence*.<sup>55</sup> For example, a realist about cat discourse is committed to the existence of cats, and to cats bearing the properties that are ordinarily attributed to them. But a corresponding realist-relevant commitment involves *independence*.<sup>56</sup> That cats exist and have the properties we attribute to them is taken by the realist to be independent of, for instance, what we believe, our conceptual schemes, linguistic practices, etc.

Alexander Miller (2012), drawing from these insights about the two key dimensions of any realism – existence and independence – offers a helpful working model of what *generic realism* (e.g. realism, for any area of discourse) is going to look like. If we let 'a', 'b' and 'c' ... be the objects that are distinctive of some subject-matter  $\alpha$ , where F-ness, G-ness, H-ness are  $\alpha$ -distinctive properties, we can characterise  $\alpha$ -realism as follows:

**$\alpha$  Realism:** a, b, and c and so on exist, and the fact that they exist and have properties such as F-ness, G-ness and H-ness is (apart from mundane empirical dependencies of the sort sometimes encountered in everyday life) independent of anyone's beliefs, linguistic practices, conceptual schemes, and so on.<sup>57</sup>

Working with Miller's simple 'two-component' picture of generic realism,<sup>58</sup> we can contrastively frame three varieties of *generic anti-realism*. The first variety denies the existence claim – that a, b, and c and so on exist.  $\alpha$ -error theorists and  $\alpha$ -non-cognitivists are, in the main,  $\alpha$ -anti-realists of this first variety.<sup>59</sup> The second variety denies the independence claim, and so allows that a, b, and c exist, but denies that their existing and the fact that they have properties like F-ness, G-ness and H-ness is a (non-trivial) mind-independent affair.  $\alpha$ -subjectivists and  $\alpha$ -idealists are typical examples of the second-kind of anti-realism. Facts about a, b and c's having F-ness, G-ness and H-ness is for the subjectivist and idealist *non-trivially* mind-dependent.<sup>60</sup> A third possible variety of anti-realist denies both the existence and independence claims.<sup>61</sup> A mad-dog solipsist will be an obvious anti-realist across the board in this third sense. We'll see how these varieties of anti-realism take shape shortly. But first, let's make things more concrete by looking at the metaepistemological case.

### 1.5.2 Metaepistemological realism

If we plug the subject matter of epistemology in for  $\alpha$ , we get a simple picture of what, *at core*, realism about the subject matter of epistemology looks like. The characteristic objects of epistemic discourse will include both individuals to whom we attribute epistemic properties (e.g. *Katie* knows there is a barn; *Leo* is intellectually virtuous), as well as distinctively epistemic objects, such as epistemic states, intellectual virtues, defeaters, etc. When we attribute properties to epistemic objects, this sometimes involves reference to *particular* epistemic objects (e.g. John's belief was rational; Stan's inquiry was intellectually virtuous). Other times, we attribute properties to abstract epistemic objects. For example, take a simple case of abstract reference: 'Knowledge requires true belief'; consider also, the universally quantified claim made by the (mentalist) evidentialist: 'for any agent, A and belief B, all facts relevant to justifying B are reflectively accessible to A'. The metaepistemological realist is going to say that when we, by way of uttering sincere declarative sentences, attribute epistemic properties (e.g. justifiedness, epistemic irresponsibility, irrationality) to things such as individuals or (abstract or particular) epistemic objects, what we've asserted will (all being well) be validated by how things stand – in a sense that involves a commitment to both existence *and* independence. Let  $a_E$ ,  $b_E$ , and  $c_E$  represent things to which we characteristically attribute epistemic properties (e.g. individuals and epistemic objects), and let  $F_E$ -ness,  $G_E$ -ness and  $H_E$ -ness be epistemic properties.

**Metaepistemological Realism:**  $a_E$ ,  $b_E$ , and  $c_E$  and so on exist, and the fact that they exist and have properties such as  $F_E$ -ness,  $G_E$ -ness and  $H_E$ -ness is (apart from mundane empirical dependencies of the sort sometimes encountered in everyday life) independent of anyone's beliefs, linguistic practices, conceptual schemes, and so on.<sup>62</sup>

The lurking critic in the wings remarks at this point with a 'gotcha' objection: knowledge involves belief, and so a claim like 'Jim knows that the bank is open' cannot be held to be true in the realist sense, because the truth of such a claim could not ever, even in principle, hold independent of anyone's beliefs, linguistic practices, etc. In fact, if the claim in question is true, it is because of Jim's belief's having certain properties (or, if Williamson (2000) is right, it's because Jim is in a particular kind of mental state). Thus, as this line goes, anyone who affirms epistemic

realism while countenancing knowledge facts embraces an internally inconsistent view.

Here it is helpful to distinguish, at least to a first approximation, between *trivial* and *non-trivial* dependencies. Consider, to this end, as Jenkins (2005) does, that: ‘One can be a realist about the physical world whilst acknowledging that there being a desk in one’s office is in some sense dependent upon on the mental states of the people who designed that desk, manufactured it, and put it in one’s office’.<sup>63</sup> Likewise, as Keller (2014) remarks, ‘Most realists are realists about minds (and mental discourse more generally), but it does not make sense to say that “there are minds” is mind-independent, or that “there are minds” would be true if there were no minds’.

The realist’s ‘independence’ thesis is thus *not* the (implausible) claim that: X’s being true does not depend *in any sense* on one’s mental life, etc. Rather, and as a working idea (we will continue to refine this as we go on), we can say that: beyond mundane dependencies, there isn’t any additional sense in which the relevant states of affairs’ being as they are depends on, as Miller puts it, ‘anyone’s linguistic practices, conceptual schemes, or whatever.’ That’s about as well as we can do for now.<sup>64</sup> It’s enough to see why the kind of dependence of facts about knowledge on beliefs aligns with the kinds of trivial, mundane dependencies that are orthogonal to the independence claim of the realist.

## 1.6 Mainstream metaepistemology and realism

With reference to the ‘two axes’ of realism – existence and independence – we can now get clearer about just how participants in paradigmatic disputes in mainstream metaepistemology reveal a tacit commitment to (at least, a *version* of) *realism*, in virtue of taking for granted, as common ground in first-order disputes, epistemic facts with an objective profile which *de facto* satisfy independence and existence conditions for realism. The picture is simple:

	Independence	Existence
Mainstream Metaepistemology	Objective	Epistemic Facts

Recall that what Moore and the sceptic were simply taking for granted, in arguing about whether K(W) is true, is that the state of the world really does validate some correct description of their claim in dispute: whether some epistemic property is instantiated by something

(e.g. Moore), or whether it is not. They both *agree* that *either* Moore knows there is an external world, or that the sceptic is right and he does not. And in this presupposed picture, thus, they are *not* taking it that there can be contradictory but *equally valid* descriptions of what the epistemic facts here are, e.g. as might be the case if epistemic facthood was just a matter of Moore's or the sceptic's opinion, or what is licensed by their respective cultural mores, and nothing more. (Recall DeRose's observation that 'each is publicly indicating that they are (or at least mean to be) contradicting the other, by saying such things as, "No, you're wrong. I do know"').

The Moorean and the sceptic are accordingly (pragmatically) presupposing not just epistemic facts *as anyone might conceive of them*, but rather their dispute takes for granted a commitment to epistemic facts with an objective profile (as it was put in §4), and as such, *not* (non-trivially) mind-dependent. (After all, as we saw in §4, *without* such objectivity presupposed, their debates would (by SPCG) take on a very different shape.) Now what is implied by the kind of objectivity that Moore and the sceptic are presupposing about epistemic facts smuggles with it in a straightforward way the realist-relevant notion of *mind-independence*, and *mutatis mutandis* for other traditional debates which in equal measure tacitly embrace a metaepistemological picture that *de facto* satisfies the *existence* and *independence* dimensions of realism about the objects and properties quantified over in first-order epistemic discourse.

I say 'at least a minimal form' because realism comes in degrees.<sup>65</sup> While paradigmatic disputes in mainstream epistemology betray a *revealed* commitment to both the existence and independence dimensions of realism, *qua* revealed commitments, there is nothing *further* that is articulated. (In fact, as I'm laying out the picture, 'mainstream metaepistemology' is not an articulated position at all, but a presupposed one.) Realism can come in more robust forms than the kind of 'minimal' realism that gets presupposed, by simply making explicit further commitments. Consider that while Moore and the sceptic reveal themselves to be taking for granted that the epistemic facts they're trying to uncover are objective in character, any *further* specifications of the metaphysical nature of such facts outstrips anything that is actually getting presupposed in the common ground. For instance, it's simply not in the common ground between Moore and the sceptic whether (if someone knows something) the properties of the belief one has are irreducible (i.e. *sui generis*) or whether they supervene on natural properties.<sup>66</sup> Mainstream epistemologists would be more than *welcome* to also endorse more 'robust' versions of realism, by positively articulating (for

instance) the irreducibility thesis, beyond what they already presuppose. But, to stress, the kind of realism they are tacitly wedded to doesn't have a horse in this further race.<sup>67</sup>

Relatedly, first-order epistemologists are also welcome to, when making explicit their metaepistemological position, *abjure* the minimal sort of realism that they are tacitly assenting to by engaging in first-order debates. This would involve making explicit why first-order epistemology proceeds with something false lurking in the common ground (e.g. that there are objective epistemic facts). Such stories are much more common in metaethics.

We are now in a position to consider how certain epistemic *anti-realist* positions will be functions of how either *existence* or *independence* would be resisted, and where epistemic relativism stands within this picture.

## 1.7 Metaepistemological anti-realism

In short, one rejects the *existence* dimension of metaepistemological realism provided that one denies – contrary to what we've seen is pragmatically presupposed in first-order epistemic discourse – that there is some fact of the matter as to whether claims such as 'John knows that *p*' or 'Knowledge must be safe' are true. *Epistemic error theory* and *epistemic non-cognitivism* both are *metaepistemological anti-realist* positions in virtue of denying precisely this. But they do so for very different reasons.<sup>68</sup> Given space constraints, and because epistemic *relativists* (who will be the focus in what follows) are *not* non-cognitivists, I will limit my focus to showing how the cognitivist *metaepistemological error theory* constitutes a straightforward form of 'anti-existence' anti-realism, and further, to note why as such mainstream epistemology proceeds as if it's false. This will be useful because clarifying just what the metaepistemological error theorist is distinguished as *denying* will help to sharpen just what is involved in denying metaepistemological realism by failing to countenance the existence of epistemic facts.

### 1.7.1 *Anti-existence metaepistemological anti-realism*

To a first approximation, the *metaepistemological error theorist* (hereafter, error theorist) rejects the existence of epistemic facts by simply denying that commonly conceived epistemic properties are ever instantiated in the world. Now, this point involves some subtlety. Consider: if one simply offers any old stipulated account of *what it would be* for an epistemic property to be instantiated in the world (and *a fortiori* any old account of what an epistemic fact would be), is that person thereby *de*

*facto* satisfying the existence dimension of realism? Surely not. After all, one might be confused and think because knowledge is beautiful, epistemic facts are just aesthetic facts. But denying the existence of aesthetic facts does not suffice for denying epistemic facts.

Here it is helpful to follow Terence Cuneo (2007) who remarks that ‘there are conceptual *limits* as to what could count as an epistemic fact. Necessarily, were epistemic facts to exist, then they would have a certain type of nature’.<sup>69</sup> Appreciating the error theorist’s claim involves appreciating these limits, which in turns involves thinking about what they might be.

Two plausible such limits which Cuneo proposes concern the *content* of epistemic facts and the *authority* of epistemic facts. More specifically, Cuneo tells us that a view countenances *epistemic* facts – as opposed to, say, some other kind of fact (e.g. moral facts) – only if countenancing facts that answer to certain commonsense *platitudes* that are characteristic of epistemic facts, and the two kinds of platitudes he has in mind concern the *content* and *authority* of epistemic facts.

Firstly, regarding *content*. Just as, Cuneo thinks, content platitudes *vis-à-vis* moral facts ‘congregate around the notion of human well-being’ (p. 58), likewise, ‘content platitudes with respect to epistemic facts cluster around the notion of *accurate representation*’. As he unpacks the idea:

...content platitudes with respect to representational entities such as the propositional attitudes display one or another epistemic merit (or positive epistemic status) such as *being a case of knowledge, being warranted, being an instance of understanding, insight or wisdom* and the like, only insofar as they are representative in some respect. That is, these entities display such merits only insofar as they *represent reality aright*. Cuneo (2007, 57)

While I think Cuneo’s line on content platitudes more or less gets it right,<sup>70</sup> it of course invites complicated questions about ‘what is necessary and sufficient’ for a fact to be an *epistemic* fact. For our purposes, we’ll just take epistemic discourse to be of the paradigmatic sort, and take (something like) the content platitude as guiding. What’s more interesting (and at stake between the error theorist and the realist) is the *authority* platitude. Cuneo (2007, 59) writes:

[Epistemic facts] are *prescriptive*... they are, imply, or indicate reasons for properly situated individuals to behave in certain ways... *regardless* of whether these agents care about conducting their behaviour in a

rational way, whether they belong to a social group of a certain kind, or whether they have entered into a social agreements with others.

For instance, as Cuneo puts it, it's implicit in the authority platitude governing epistemic facts that if it's a fact that I have an epistemic reason to believe some proposition for which evidence is compelling, then *that I have this epistemic reason* is 'not contingent on whether I care about believing what is true' (*ibid.*, 59). Cuneo is certainly right that epistemic reasons for belief are not construed in epistemology as 'hypothetical' reasons; and the commonsense conception we have of epistemic facts, more generally, aligns with what we say about epistemic reasons. Put simply, *whatever the epistemic facts are*, they hold in a way that is *not* (non-trivially) dependent on our having certain desires, occupying social roles, etc. Respecting conceptual limits as to what could count as an epistemic fact, thus, must involve respecting *this point*.

The foregoing suggests what is involved in taking the error-theoretic route to epistemic anti-realism; in short, the error theorist denies the *existence* leg of metaepistemological realism, and this is because the error theorist denies that there are any epistemic facts that satisfy the kinds of commonsense platitudes characteristic of epistemic facts, *viz.* – that no *commonly conceived epistemic properties are actually instantiated in the world*. And accordingly, because there are no such facts whereby such properties are instantiated, the error theorist's assessment is that claims to epistemic fact are categorically false. To my knowledge, there are no practicing mainstream epistemologists who are epistemic error theorists. In fact, the only contemporary epistemic error theorist I'm aware of is Jonas Olson (2009) – and Olson is first and foremost a metaethicist and value theorist, one whose philosophical rationale for epistemic error theory is premised upon a more general kind position grounded in concerns about moral properties and moral facts.

Regardless of what we say about *moral* error theory, it remains that if *epistemic* error theory were true, it would be entirely puzzling why anyone would try (with any seriousness) to do any *first-order epistemology*. One way to appreciate this point is to reposition Crispin Wright's (1992, 9) criticism of Mackie's moral error theory as a criticism against the epistemic error theorist, specifically. The reimagined Wright-style criticism goes like this:

The great discomfort with [epistemic error theory] is that, unless more is said, it simply relegates [epistemic] discourse to *bad faith*.

Whatever we may once have thought, as soon as philosophy has taught us that the world is unsuited to confer truth on any of our claims about what is [epistemically] right, or wrong, or obligatory, etc., the reasonable response ought surely to be to forgo the right to making any such claims... If it is of the essence of [epistemic] judgement to aim at truth, and if philosophy teaches us that there is no [epistemic] truth to hit, how are we supposed to take ourselves seriously in thinking the way we do about any issue which we regard of major [epistemological] importance? (Wright (1992, 9), *my italics and with modifications*)<sup>71</sup>

This is not the place to attempt to argue that epistemic error theory is false.<sup>72</sup> It will suffice to consider why – for reasons akin to the Wright line – just as soon as one begins engaging in debates that actually *occupy* first-order epistemologists, one is proceeding, in practice, as if epistemic error theory is false, as if there *are* epistemological facts to uncover.

### 1.7.2 Anti-independence metaepistemological anti-realism

I want to turn now to a second and very much related road to metaepistemological anti-realism, which comes *via* a denial of the *independence* leg of realism. Denying the independence leg, recall, involves the view that the states of affairs (e.g. consisting in the instantiation of epistemic properties in the world) that make up the epistemic facts are *non-trivially* mind-dependent.

As we saw in §1.5, it's not clear that there is to be found any general formula for distinguishing trivial from non-trivial dependencies, such that the general formula can be applied universally across subject matters. We did point out several instances of obvious *trivial* dependencies (e.g. table-facts depend, in a trivial sense, on the mental states of their makers. If, as Williamson (2000) thinks, knowledge is a mental state, then knowledge-facts obviously depend on the mind, but in a way that is trivial). Let's now consider, in the arena of epistemic facts, an example of a candidate *non-trivial* dependency claim.

Let's go extreme. Here I have in mind, *a la* Berkeley, a view we can call *metaepistemological idealism about epistemic facts*, a toy view about epistemic facts according to which the instantiation of epistemic properties just *is* an idea in the mind,<sup>73</sup> given that reality *itself* is regarded as mental in nature; on this view, the states of affairs that obtain when one knows a given proposition will necessarily be mentally constituted – and thus, *anything we know* we know because some mental state of affairs obtains. Obviously, on such a view, it would seem initially

obvious why the ‘non-triviality’ aspect of denying mind-independence would be met – after all, this proposal renders epistemic facthood about as mind-*dependent* as you can get! And since the envisaged Berkeleyan metaepistemological idealist about epistemic facts is neither an error theorist nor a non-cognitivist about epistemic facts – but rather, she (stipulate) tells us there are epistemic facts – the natural way to read this toy view is as follows: as one that embraces (as self-advertised) the existence of, albeit peculiar, epistemic facts but which is nonetheless anti-realist along the independence axis – because the epistemic facts embraced are non-trivially mind-dependent. At least, that looks like the official story.

But things here are a bit more delicate than they would seem, and just *why* this is will be instructive for how we should think about the complex relationship between mainstream metaepistemological realism and relativism. To ease into what will become an important point as we go on, consider Sayre-McCord’s (1991, 158) remarks:

[Berkeley] offers his version of idealism as a literal construal of material object statements and so saw himself as defending realism – yet his construal is so implausible that his position is almost universally regarded (rightly) as an anti-realist position, *despite his protestations to the contrary*. (p. 158, my italics)

*Protestations to the contrary?* Question: does a Berkeleyan metaepistemological idealist about epistemic facts have any *remotely* plausible cause to fancy herself a realist, if we understand (as I’ve suggested we do) realism along the existence and independence axes?

The answer is, awkwardly, yes. After all, the Berkeleyan will tell us there *are* epistemic facts, albeit, facts which hold in virtue of mental properties being instantiated by mentally constituted entities. *But then* suppose she proceeds to tell us – and at this point, we are a bit stunned – that she regards these epistemic facts as merely *trivially* mind-dependent. After all, the Berkeleyan points out, *if material idealism is correct* and the whole world is mentally constituted, then it *trivially* follows that representing aright *could only be* representing mentally constituted states of affairs.

To a first approximation, here’s why the envisioned Berkeleyan metaepistemological idealist about epistemic facts should ultimately fail to convince us she’s a realist by adverting to the line just considered: it’s because *the way* that Berkeleyan metaepistemological facts are mind-dependent will have, as a knock-on effect, that the *existence* leg of realism

isn't going to be plausibly satisfied. And this is so *regardless* of whether the Berkeleyan claims she countenances epistemic facts.

Crucially, *even if* one purports to embrace the existence of epistemic facts but in doing so opts for a characterisation of these facts on which epistemic facthood is, say, a radically mind-dependent affair, then *what one's got left isn't necessarily going to be epistemic facts anymore*. And this is for the simple reason that: (i) *a la* Cuneo (2007), there are conceptual *limits* to what counts as an epistemic fact – *viz.*, such facts must satisfy the content and authority platitudes; and (ii) one's *characterisation* of epistemic facts, and how they should be understood as depending on our mental lives, can itself be a difference-maker with respect to whether the would-be epistemic facts satisfy content and authority platitudes.

To appreciate this point, we need only consider why the *kind* of 'epistemic facts' our envisioned Berkeleyan idealist countenances (obviously) fail to be *authoritative* in the way that epistemic facts, as such, must *a la* Cuneo, *necessarily* be. Suppose A and B are intrinsic duplicates with identical microphysical causal histories. Even in such a case, it will come out false on the envisioned Berkeleyan view that any evidence E possessed by A will be prescriptive for B. After all, A's having evidence E *vis-à-vis* some proposition *p* which is *about* how things stand in A's *A-mentally-constituted* world would not be authoritative in any way *vis-à-vis* how B conducts her cognitive affairs in B's *B-mentally-constituted* world.

But there's an even deeper problem about authority here. It's hard to see how (on the toy Berkeleyan view) some evidence E possessed by A could ever be *authoritative for A*. The reason is that the prescriptive force of epistemic facts as such (according to the authority platitude) is necessarily *non-hypothetical* – a point a Cuneo-style realist as well as an Olson-style error theorist agree upon. However, given that perception grounds essence on the (very strange) view under consideration, one would be in a position to nullify any would-be epistemic reason-giving force E has for A by simply perceiving *differently*. Put more abstractly: on the Berkeleyan view, epistemic reasons are always going to be wide scoping over perception → essence conditionals, where the giving up of the antecedent (e.g. the changing of the perception) cancels the reason.

The take-away point from thinking about the – admittedly, bizarre – example of the Berkeleyan metaepistemological idealist is this: even though the Berkeleyan (unlike the error theorist and noncognitivist) *claims* to countenance epistemic facts, the particular *way* the Berkeleyan opts to render epistemic facts as mind-dependent relegates (and *regardless* of whether the mind-dependence is ultimately regarded as trivial or non-trivial) the epistemic facts the Berkeleyan preserves as *failing* to

fall within the conceptual limits of what counts as an epistemic fact, *in virtue of* being the sort of things that fail to satisfy the authority platitude. And accordingly, the envisioned Berkeleyan *de facto* fails the existence dimension of realism, despite whatever her protestations may be, and regardless of whether she insists that the mind-dependence of epistemic facts, *conditioned* upon the truth of her view, is trivial.

And from this we can generalise: given that there are plausible conceptual limits to what count as epistemic facts, *any* metaepistemological thesis is going to be anti-realist along the *existence* axis provided that this thesis allows for mind-dependence in a manner such that the only 'epistemic facts' left standing are ones that fail to satisfy either content or authority platitudes.

### 1.8 Metaepistemological realism and relativism: an impasse

Against this background, we're now in a position to approach the very complicated issue of where the *epistemic relativist* stands in connection to the kind of metaepistemological realism revealed to be taken for granted in typical first-order debates in mainstream epistemology. Unlike metaepistemological idealism – which is grounded in a more general thesis, material idealism, that is more or less target practice – *epistemic relativism* has had a long and complicated relationship within (and from the periphery of) the theory of knowledge. While there are a variety of different versions of epistemic relativism – versions we'll examine in detail in the coming chapters – let's consider here a very simple version of the view, and then consider how this view lines up with mainstream metaepistemological realism.

To make things as straightforward as possible, consider a (simplicification of) a version of epistemic relativism often attributed to Richard Rorty (1980). Call this simple view *cultural relativism about epistemic justification*. According to this proposal, what is being denied is not that there is a fact of the matter whether (for instance) in a given cultural context *C*, *S* will be justified in believing some proposition *p*. Like the Berkeleyan, and unlike the error theorist and non-cognitivist, the cultural relativist about epistemic justification tells us that there are epistemic justification facts and that our claims to represent these facts are truth-apt. However, the Rorty-style relativist (also, like the Berkeleyan) has a particular idiosyncratic story about the character of such facts.

On the cultural relativist's story, there will be (to put it crudely for now) *perspective-relative facts* about (for example) what epistemically

justifies what.<sup>74</sup> In the simple case where the state of affairs that would ordinarily suffice as the ‘epistemic fact’ (e.g. the state of affairs of an agent instantiating an epistemic property, say, the property of being justified in believing that *p*), the relativist conceives of the state of affairs that constitutes the relevant (perspective-relative) fact as the state of affairs where *some agent instantiates an epistemic property E-according-to-some-perspective-p*.<sup>75</sup> When, for example, the relevant perspective is the perspective of Western science, the Rortian cultural relativist about epistemic justification will tell us that the claim ‘The doctor’s belief is justified’ is true, and that there is a fact of the matter; the fact is, however, a state of affairs where the doctor instantiates an epistemic property *according to the perspective of Western science*. There are no perspective-independent epistemic facts. If we suppose there is another perspective, the perspective of Astrology, then relative to that perspective, it might be that the doctor’s belief is epistemically unjustified; the relativist denies there is any *further* perspective-independent fact of the matter, about which perspective is ‘right’.

To reiterate, epistemic relativism comes in (much) more complicated versions than the one just outlined – several of which we’ll consider in further chapters – but this simplification of Rorty’s view will do for now, in so far as I want to close this chapter by setting up a kind of ‘cliff-hanger’ to do with the relationship between mainstream epistemology and epistemic relativism.

The cliffhanger gets off the ground once we pose a question: is the (simple) epistemic relativist picture just sketched an example of metaepistemological realism, or not? This answer to this question will be important, because if *not*, then epistemic relativism is, as a metaepistemological position about epistemic facts, simply *incompatible* with even the very general *kind* of metaepistemological view which, as I’ve argued, mainstream epistemologists take for granted to get their projects off the ground – *viz.*, a kind of realism. And if this is right, then retrospectively, we should not be surprised that epistemic relativism has been more or less banished to the outskirts of mainstream epistemology, even if it remains popular elsewhere.<sup>76</sup> Even more, if this metaepistemological incompatibility can be established, then it looks like there is a positive reason why mainstream epistemologists in fact *don’t* take the epistemic relativist seriously: in short because the presuppositions of their first-order projects tacitly exclude it.

Does this ‘incompatibility question’ (as I’ll refer to it later, e.g., in Chapter 9) get an affirmative answer? This is complicated; in short, *it depends*. Consider that it might well be that the perspective-relative facts

the epistemic relativist countenances under the banner of ‘epistemic facts’ fall *outside* plausible conceptual limits of epistemic facthood, by failing to satisfy the content and authority platitudes. Let’s grant for the sake of argument that perspective-relative epistemic facts satisfy Cuneo’s *content* platitude. It is the remaining matter of whether perspective-relative epistemic facts satisfy the *authority* platitude that is, I think, where the action lies. And by action, I mean: here we reach a very awkward kind of impasse – a cliffhanger.

Consider that *if the Cuneo-style (non-relative) realist is right*, then it looks like the epistemic relativist is *not* going to qualify as a metaepistemological realist, and this because the epistemic relativist fails to satisfy the Cuneo-style-realist’s idea of what it takes to satisfy the authority platitude according to which the notion of authority should be unpacked *inter-perspectivally*. That is: if we have more reason to embrace a Cuneo-style (non-relative) realism about epistemic facts than we do to embrace a relativist picture of epistemic facts, then we also have a reason to insist on *a certain conception of what is required to satisfy the authority platitude*, one according to which *intra-perspectival authority* (of the sort that characterises what would be the epistemic relativist’s claimed epistemic facts – e.g. that they are authoritative *within* but not *outside* a particular perspective) is *not* sufficient.

*However*, the other side of the coin is that if we have good reason to embrace the wider picture of epistemic facthood the relativist is selling – i.e. *if arguments for epistemic relativism turn out to be compelling ones* – then it looks like the epistemic relativist is going to be best classed as a *metaepistemological realist*, and this because the epistemic relativist will straightforwardly satisfy the *authority platitude unpacked as it must be if epistemic relativism is true* – *viz.*, that epistemic facts have an authority that transcends desires and goals (i.e. an authority that is not merely hypothetical, of the sort Olson can accept) but *not epistemic frameworks/perspectives*. And so: if we have compelling reason to embrace the epistemic *relativist’s* picture of epistemic facts, then we have at the same time reason to insist on a certain conception of what is required to satisfy the authority platitude, one according to which *intra-perspectival authority* is sufficient.

It should now be evident why the very *plausibility of epistemic relativism* is ultimately crucial to how we should answer the question of incompatibility, no less than the more general question of whether mainstream epistemology’s lack of engagement with the epistemic relativist is with or without good philosophical basis – an issue I’ll revisit in much more detail in Chapter 9. In short, if there is good reason to accept (at least

some version of) epistemic relativism and the wider picture of epistemic facts that falls out of this view, then the epistemic relativist can't simply be written off as not *compatible* with mainstream metaepistemological realism on the grounds that the epistemic relativist is not a realist. *But* if epistemic relativism is *not* plausible, then, *a fortiori*, *neither* is the picture of the conceptual limits of epistemic facthood it recommends and with reference to which the epistemic relativist could be regarded as a full-blooded realist. If epistemic relativism is not plausible, then (interestingly) we lack any reason to think epistemic relativism is a version of epistemic realism – and correspondingly, have cause to think the view is simply *incompatible* with the kind of view that mainstream metaepistemology takes for granted. Thus, the road so far is one that leads to one straightforward question: *how plausible is epistemic relativism, really?*

Answering this involves engaging with epistemic relativism of many different stripes. In the next chapter, we'll look at the most radical path to epistemic relativism – one which proceeds through *global* relativism.

# 2

## Global Relativism

*Abstract.* If *everything* is relative, then so are epistemic standards, norms and facts; is there anything to recommend the quick path to epistemic relativism, *via* global relativism? This chapter explores this question by examining global relativism in some detail. Of particular interest will be the familiar charge that the doctrine is self-refuting. Careful attention to the way Plato attempted to level this very charge against Protagoras's version of global relativism will be instructive, and it will be shown how Burnyeat's (1976) nuanced defence of Plato's self-refutation argument in the *Theatetus* foreshadows an important dividing line in contemporary thinking about relativism between Boghossian (2006a) and Wright (2008). The conclusion reached in this chapter is that global relativism is ultimately not a defensible view, though this is hardly because (as is often thought) it can be dismissed with a quick 'knock-down'.

### 2.1 Introduction

Global relativism is, put crudely for now, the thesis that '*everything* is relative'. One rationale for epistemic relativism proceeds like this: if global relativism is true, then so is epistemic relativism. Global relativism is true. So epistemic relativism is true, too. Call this the *trivial entailment* argument for epistemic relativism.

Is this argument any good? Granted, hardly anyone in professional philosophy self-identifies as a global relativist. If asked about the position, a very common response is to quickly dismiss the view out of hand as something like, as Boghossian (2006a, 54) puts it, '*a fundamentally*

incoherent position' – and perhaps to then suggest moving on to something worth taking seriously!

If global relativism could be so easily rationally dismissed as you'd expect, given how quickly it in fact *is* brushed aside, then the thought that we should take seriously the suggestion that epistemic relativism might be motivated off the back of global relativism looks at best under-motivated. But global relativism is, in fact, surprisingly hard to dismiss. It might be a *prima facie* outlandish target, but even if that's right, it remains that global relativism is a difficult and slippery target to actually hit.

One reason global relativism is a slippery target is that it's not very easy to picture what things would even *be* like if the thesis were true.<sup>1</sup> This, as we shall see, is an issue that raises its head in a number of key places where friends and foes of the thesis have famously clashed – as does the connected, and complicated, issue of how, precisely, we should think about the relativised truth predicate that the global relativist either explicitly (or implicitly) adverts to, and this is especially so in the special case where the doctrine is applied to itself.

Yet a further complication for evaluating global relativism concerns what has been, both in ancient and contemporary discussions, the principal philosophical objection to the thesis, which is that it is *self-refuting* – *viz.*, the charge that, as Kölbel (2011, 11) puts it, the thesis can 'in some way be turned *against* itself' (though there are different ways this could happen). This chapter examines in some detail how such a charge might actually be seen to put the global relativist to rest, and *a fortiori*, the prospects of a trivial entailment argument for *epistemic* relativism. It turns out that putting the global relativist to rest is every bit as tricky as one, on reflection, might fear.

## 2.2 Relativism itself

A few preliminaries about relativism are in order before locking horns with its global version. The first is definitional. What is *relativism*, exactly?<sup>2</sup> There are broadly three ways this question has traditionally been addressed,<sup>3</sup> and we can capture the gist of all three by using a simple example of a paradigmatic relativist thesis: relativism about *moral values*.

Consider the claim:

(MV): Moral values are relative to local cultural norms.

One way to think about why MV is a relativist thesis is negative – *viz.*, in terms of what a proponent of MV denies. Rorty (1979) and Derrida

(1974), for instance, are often branded relativists on the basis of the kinds of theses they are noted for rejecting<sup>4</sup> about some class of truths under consideration; for instance, relativists characteristically reject theses that go under the description of *absolutism*, *objectivism* and *monism*, in virtue of maintaining that (*contra* absolutism) the truths under consideration are not universal – *viz.*, applicable to all times and frameworks; that they are (*contra* objectivism) not independent of our judgments and beliefs; and (*contra* monism) that that competing viewpoints need not mutually exclude one another.<sup>5</sup> We might say, then, that MV is a relativist view in virtue of denying these three interconnected theses about moral truths.

A second approach to characterising relativism is the *co-variance* approach (e.g. Baghranian 2004; 2014 and Swoyer 2014), which thinks of relativism in terms of the relationship between two variable places that are useful for classifying different *kinds* of relativism – *viz.*, in terms of what is relativised (the object of relativisation), and what it is relativised *to* (the domain of relativisation). On the co-variance view, a thesis (e.g. MV) will count as a version of relativism provided it submits that some object, *x*, depends on some underlying, independent variable, *y*, such that, in some suitably specified sense, change in the latter results in a change in the former. Therefore, what makes MV *relativistic* on this approach is that MV submits a view where local cultural norm change instances moral value change.<sup>6</sup>

While the co-variance approach is useful for taxonomising relativisms, one might wonder whether, for the purpose of distinguishing the intuitively relativist from non-relativist views, it is too *inclusive*, and that this problem will crop up in principle no matter how the relevant change relation is fleshed out. Consider human beings, on the one hand, and the earth's oxygen levels, on the other. Human beings *depend* on oxygen; a change on oxygen levels on earth will result in very drastic changes in human beings. As the worry goes: everyone accepts this; but this isn't sufficient for making everyone *relativists about human beings* in the kind of provocative sense that we take a claim of relativism proper (like MV) to involve. And *mutatis mutandis* for tighter relationships – e.g. between weight measured and the kind of scales used.

One response to this inclusiveness point is to simply note that the philosophically *interesting* forms of relativism will be ones where the relevant dependence is invoked to settle (or explain away) what look like profound disagreements, where the relativising parameter often involves the kinds of things that are initially *responsible* for the disagreements in the first place (belief systems, cultures, conceptual frameworks, epistemic frameworks, languages, etc.).<sup>7</sup> On this rationale, MV

is a philosophically interesting kind of relativism then (unlike the case of humans and oxygen, or weight and scales)<sup>8</sup> on the co-variance approach because moral values, characteristically a point of profound disagreement, are on MV being relativised to local cultural norms, where differences in the latter are among the salient contributors to such disagreements.

A third approach to articulating relativism is what we can call, following Spencer (2014), the *arity* approach, on which relativism is understood in terms of properties and their *degree of relationality* – viz., their *arity*. A central proponent of this characterisation of the core relativist insight is Crispin Wright (2008a).

The ground-level relativistic idea is that the satisfaction-conditions of a certain property or family of properties, *though superficially presenting as unary*, are actually implicitly relational – or more generally, are of a higher degree of relationality than is apparent in the surface syntax... [The] tacit relationality need not be to the effect that a certain apparently unary property is in fact binary. It may be to the effect that a certain apparently  $n$ -ary property is in fact  $n + k$ -ary,  $k > 0$ .<sup>9</sup>

The condensed definition of relativity on the arity approach is, as Spencer formulates it:

*(Arity) Relativity*: an apparently  $n$ -ary property is relative to a parameter only if the property is  $n + k$ -ary,  $k > 0$ .<sup>10</sup>

What is an *apparently  $n$ -ary* property? On the Wright line, a property is *apparently  $n$ -ary* if (given surface syntax) the property ‘superficially presents’ itself as of the degree of relationality  $n$ . And this turns out to be an important idea; the corollary to the ‘appearance’ condition is, of course, that the property in question is claimed to be *not* of the *apparent* degree of relationality, but rather of a higher degree – e.g. not ‘beautiful’ (*simpliciter*) but beautiful in relation to some parameter (e.g. some perspective).

Gilbert Harman (1975), Robert Nozick (2001) and Paul Boghossian (2006b) have all, in different ways, tried to capture the thrust of this relationality insight when defining relativism in terms of some parameter’s being *hidden* or *unexpected* in light of how things appear, given some specification of a background consisting in facts about how we talk, what we know, etc.<sup>11</sup> MV is a form of relativism on the arity approach

because (surprise!) MV commits us to saying that (for example) an act  $\phi$  is never *just* wrong, despite our thought and talk seeming to vindicate that some things just wrong; rather on MV  $\phi$  is *wrong* always only once some further parameter is specified, and in the case of MV, that parameter is supplied by local cultural norms. The matter of how to *think* about that parameter and the role it plays with respect to the truth in question is a matter of deep divide among those endorsing an arity approach to relativity, and I'll return to this at a later juncture.

But first, a more general point is in order. Consider that thinking about relativism in terms of arity (and, in particular, by invoking talk of 'additional parameters') is going to be a natural strategy for philosophers interested in the semantic dimensions of relativism.<sup>12</sup> An arity relativist who says that  $\phi$  is wrong always only once another parameter is specified (and never wrong *simpliciter*) might well consider, as a *restatement* of this view, the claim that 'x is wrong' is *true* (or more generally, gets a truth value) always only once some parameter is specified. In fact, Paul Boghossian, who embraces a version of the arity approach to thinking about relativism (as involving always an unexpected parameter) thinks of relativism in terms of *arity-relativity*, and arity-relativism as just equivalent to a truth-referencing articulation of the relativist's insight; he writes: 'the relativist about a given domain, D, purports to have discovered that the *truths* of D involve an unexpected relation to a parameter' (Boghossian (2006b, 13), *my italics*).

The general point intimated by talk of relativism in terms of degrees of relationality is this (note: a point not *exclusive* to arity-approaches to defining relativism): that relativism is inextricably an *alethic* thesis.<sup>13</sup> Relativism *about* some feature, *F*, is very plausibly expressible *in terms of* relativism about truth.<sup>14</sup> And this general point should not be surprising. After all, as Kölbel (2011, 18) puts it in his discussion of the connection between relativism about some feature and relativism about truth, 'truth is conceptually connected to other features, and this creates also a conceptual connection between relativism about truth and relativism about other features'.

Accordingly, and provided we embrace the equivalence schema that: it is true that *p* if and only if *p*, 'relativism about any feature entails relativism about truth, and relativism about truth entails relativism about some feature'.<sup>15</sup> We'll return to this general point as it crops up (albeit, somewhat complicatedly) in the special case of global relativism.

This point about relativism and truth in hand, let's consider one more ground-clearing distinction not yet articulated: that between local and global relativism. Most relativisms which philosophers actually *defend*

are *local* in the sense that what they apply to is limited to some domain.<sup>16</sup> Local relativisms are, modulo certain entailments, theories *about* what they say they're about. Typically, local relativisms are endorsed on the basis of philosophical reasons connected to the kinds of *features* that are claimed to be relative (e.g. aesthetic standards, epistemic principles), or relatedly, semantic considerations to do with discourse where such features are attributed.

But you might endorse relativism not for any *particular* area of discourse; you might just think – for whatever reason – that all *truth* is relative – *viz.*, and with reference to the equivalence schema, this is part and parcel with the position that it is always *relative* (e.g.  $n + k$  ary,  $k > 0$ ) whether or not any feature claimed to apply to something does so apply. This is more or less the idea attributed to the global relativist, that *everything is relative*.

With reference to the co-variance approach, we can attempt to get a further grip on the thought that everything is relative by asking: relative to *what*? And here it's incumbent on the global relativist to specify some parameter, or domain of relativisation.<sup>17</sup> Protagoras, as we'll shortly see, had a simple answer to the domain question: man.

### 2.3 Protagorean global relativism: a case study

Protagorean global relativism, and in particular, Plato's controversial attempt to refute this doctrine in the *Theatetus*, has really been (given the tremendous body of classic and contemporary analysis devoted to it) the *locus classicus* in philosophical analysis of global relativism, much as Descartes' *Meditations* is the *locus classicus* in the philosophy of global *scepticism*.

Engaging with this case study will be especially helpful for appreciating just how, exactly, a plausible refutation of the global relativist, replete with counters to her most natural available escape routes, might best be thought of as proceeding. Plus, by digging in to Plato's famous *peritrope* argument in some detail, two fundamental (and recurring) insights about relative truth are brought into sharper relief. Firstly, an appreciation of the subtleties of Plato's argument against Protagoras, in particular as they have been drawn out by Burnyeat (1976), illuminate some of the important, and essential, anti-realist elements of global relativism – ones which sometimes fly under the radar – and which will bear on our later discussion of local varieties of relativism, such as epistemic relativism. Secondly, we'll see that what is perhaps the best way to make sense of Plato's attempted refutation of Protagoras in the *Theatetus* involves a

crucial appeal to what Burnyeat calls the ‘translation principle’, which is a principle that turns out to be at the very heart of a more contemporary debate about relativism (gestured to earlier, in the discussion of relativism and ways of unpacking *arīty*) between Boghossian (2006a; 2006b), Kusch (2010) and Wright (2008).

### 2.3.1 The peritrope: a first pass

Protagoras famously defended the ‘Measure Doctrine’, the doctrine that (in crude form) man is the measure of all things. What does this mean? Sextus Empiricus mistakenly thought the Protagorean measure doctrine boiled down to the radical subjectivist thesis that every *appearance* whatsoever is true (*simpliciter*). And radical subjectivism so construed quickly unravels: it can *appear* that the thesis that ‘every appearance whatsoever is true’ is false. So from the measure doctrine-cum-radical-subjectivism, the measure doctrine is false. At any rate, Sextus attacked Protagoras’s doctrine along precisely these lines, insisting that it is *self-refuting*. But no matter how implausible Protagoras’s view actually was, the thesis that *every appearance is true* is a really a strawman version of it.

The key line Protagoras actually defended, in espousing the measure doctrine, was that every judgment is true *for* (in relation to) the person whose judgment it is.<sup>18</sup> Now, given that this is the kernel of the view, it is, as will shortly be apparent, very awkward to consider how Plato himself famously claimed, in the *Theatetus*, that Protagorean global relativism is self-refuting.<sup>19</sup> Consider here Plato’s most famous line of reasoning against Protagoras, an argument referred to under the description of the *peritrope* (περιτροπή) or ‘self-refutation’ argument:<sup>20</sup>

...it [The Truth that Protagoras wrote] has this exquisite feature: on the subject of his own view, agreeing that everyone judges what is so, he for his part [Protagoras] presumably concedes to be true the opinion of those who judge the opposite to him in that they think that he is mistaken... Accordingly, he would concede that his own view is false, if he agrees that the opinion of those who think him mistaken is true. (*Theatetus*, 171ab)

In what sense is Plato alleging that Protagoras’s doctrine is self-refuting? The quick and ready answer seems to be that Plato thinks Protagoras’s own doctrine ends up in a kind of unavoidable *contradiction*. But here we have to be careful, and it is helpful to distinguish, following Kölbel, two subtly different ways a doctrine might be self-refuting: *contradictorily* and *dialectically*. As Kölbel notes, contradictory self-refutation takes (at least)

two distinct forms. Call a thesis *strongly* contradictorily self-refuting if the content of the thesis entails the falsehood of the thesis – *viz.*, perhaps<sup>21</sup> the liar sentence ‘This statement is false’; a thesis is by contrast *weakly* contradictorily self-refuting if the statement in conjunction with some *a priori* premise entails the falsity of the thesis. This would be what we find in a case where an individual (e.g. Socrates) claims to know that he knows nothing. This statement in conjunction with the *a priori* premise that what is known is true jointly entail that Socrates does not know anything and then (contrary what Socrates has just said) he does not know that he knows nothing.<sup>22</sup>

It’s not at all clear that the quoted passage is attempting to show that Protagorean relativism is contradictorily self-refuting, in either the weak or the strong sense. After all, the measure doctrine does not, like the liar sentence does, entail its own falsity without some further premise. However, the premise adverted to in the passage at [171ab] of the *Theatetus* – that Protagoras’s opponent judges his doctrine to be false – is not an *a priori* premise, as was (in the example of the weakly contradictorily self-refuting statement that one knows one doesn’t know anything) the premise that what is known is true. Rather, the premise that some of Protagoras’s opponents regard his doctrine as false is just an *obvious* premise, albeit one that Plato thinks is not up for debate whether Protagoras accepts it (and what Plato thinks that by Protagoras’s own lights would follow).

Here I think Kölbel is right to regard the line of reasoning in the above passage to reflect that the kind of self-refutation Plato was aiming for with Protagoras is what he calls *dialectical* self-refutation, though an explanation for what this involves is subtle. A thesis is *dialectically self-refuting* when, under certain rules agreed to, explicitly or implicitly, by both parties,<sup>23</sup> the proponent of the thesis in question is forced (by her own lights) into contradicting herself. One such very minimal rule is that you must acknowledge your interlocutor’s statement. Plato seems to suggest that once Protagoras acknowledges, as he must, that at least some opponent regards the measure doctrine to be false; then by the lights of the measure doctrine, the measure doctrine is false – and in such a way that Protagoras must concede this.

But a straightforward worry immediately surfaces. Protagoras’s assertion of his measure doctrine, (M) would, in conjunction with the premise that someone judges that M is false, seem to force Protagoras into contradiction only if the ‘is true’ predicate operant in Protagoras’s measure doctrine, M, is understood as unqualified – *viz.*, as the doctrine that every judgment (as opposed to appearance) is true (period), as opposed to *true for the person whose judgment it is*. After all, it would seem

that only if we are operating with an unrelativised truth predicate will the falsity of M for Protagoras's opponent imply that it must also be false for Protagoras himself (who judges it true).

But Protagoras's doctrine does *not* say that if someone judges that  $x$ , then  $x$  is true (full stop). And so from the measure doctrine, in conjunction with the premise he is obliged to grant – that his opponent judges M to be false – Protagoras is really only left in a spot where he must admit the following, as Burnyeat (1976) calls it, 'interim conclusion':

*Interim Conclusion (IC):* The measure doctrine is false for Protagoras's opponents.

And, being forced to admit IC *appears* to be completely harmless to Protagoras, who is (given M) in a position to rebut something to the effect of: 'My doctrine might be false for my opponents, but that doesn't mean it's false *for me*'. This might sound annoying; however it's certainly the natural move Protagoras would seem *obligated* to make. And it's not a *concession* at all, and nothing that, in any obvious way, should be thought by Protagoras's own lights objectionable.

At any rate, this is all very awkward. It *looks* as though Plato has mucked things up twice over: that he fails to show that Protagorean relativism is self-refuting and further that he didn't really *understand* Protagorean relativism<sup>24</sup> (and indeed, that the failure to understand explains why his argument is ineffectual). But perhaps things here aren't as they seem.

Burnyeat (1976) has offered a very different way of looking at things – on which Plato's peritrope argument can be appreciated as, and despite the initial diagnosis just suggested, solidly establishing that Protagorean relativism *really is* (dialectically) self-refuting, and on Burnyeat's interpretation (and contrary to how things seems) Plato knew what he was doing all along.<sup>25</sup> For our purposes, it matters less whether Plato knew what he was doing – a matter for Plato scholars to debate – and more whether Protagorean global relativism is self-refuting (and if so how it might be *shown* to be so).

### 2.3.2 Formulating PGR

Without a doubt the 'major contribution' of the view Protagoras is submitting *via* his measure doctrine is captured in his statement that 'what seems to each person is so for the person for whom it seems' (170a). Call this component of the view *no-mistakes*.

*No mistakes (NM):* If  $S$  believes that  $p$ , then it is true for  $S$  that  $p$ .<sup>26</sup>

Importantly, though, Protagorean global relativism is not *exhausted* by (NM).<sup>27</sup> As Burnyeat observes, it is clear that Plato also understands Protagoras to hold *the converse*,<sup>28</sup> according to which a proposition is true for an individual *only if that individual believes that proposition*. Call the converse *no-unbelieved-truths*:

*No Unbelieved Truths* (NUT): If it is true for *S* that *p*, then *S* believes that *p*.

*No-mistakes* and *no-unbelieved-truths* are implied by the full expression of the ‘measure’ doctrine: ‘man (sc. each man) is the measure not only of what is (sc. for him), but also of what is not (sc. for him)’ (*ibid.* 178). Put another way, Protagorean global relativism boils down to the biconditional that follows from the conjunction of *no-mistakes* and *no-unbelieved-truths*:

*Protagorean Global Relativism* (PGR): *S* believes that *p* if, and only if, it is true for *S* that *p*.<sup>29</sup>

Interestingly, the fact that PGR’s full expression involves *both* of these claims means that *even if* a partial version of the thesis – PGR-cum-no-mistakes – isn’t subject to some malignant version of a self-refutation charge, it wouldn’t thereby follow that PGR (in all its glory) isn’t subject to a self-refutation charge. After all, it might be that PGR, in virtue of both NM *and* NUT, is self-refuting. And in fact that’s what Burnyeat is convinced Plato thought. Though as we’ll see, NUT– *viz.*, the no-unbelieved-truths element of the PGR, is a bit mysterious, as is the related matter of how to delineate the *consequences* of embracing the biconditional which NUT’s inclusion in the formulation of PGR generates.

### 2.3.3 Synonymy, anti-realism and the ‘secret doctrine’: a dilemma

Let’s take a closer look now at the biconditional characterising PGR – *S* believes that *p* if, and only if, it is true for *S* that *p* – and how the force of the dialectic against Protagoras is sensitive to the fact that he’s committed this way, and not merely to the left-to-right direction (i.e. NM). To make this idea more concrete, let *S* be Socrates and *p* be some proposition, never mind what (for now!).<sup>30</sup> The biconditional Protagoras is committed to defending implies (in our concrete case) the following: that (by NM) *if Socrates believes that p, then it is true for Socrates that p* and by NUT that *if Socrates does not believe that p, then it is not true for Socrates that p*. So far so good.

Now, along with the biconditional, Protagoras is committed (by, *a la* Kölbel, the dialectical rule of acknowledging your interlocutor's claim) to acknowledging the contingent fact that Socrates for his part does *not* believe that anything *he* (Socrates) believes is *thereby* true<sup>31</sup> – that is, that Socrates does *not* believe what is implied for him by NM, namely, that *if he, Socrates, believes that p, then it is true for Socrates that p*.

And by acknowledging this, Protagoras must accordingly accept (by implication from NUT) the following claim, which is a modified version of the claim that earlier appeared innocent under the description of 'interim conclusion'.

*Interim Conclusion\** (IC\*): *It is not true for Socrates that: if Socrates believes that p, then it is true for Socrates that p.*<sup>32</sup>

Protagoras is, by his own doctrine and what he is dialectically obligated to acknowledge, committed to IC\*. Now, recall that (originally, at least) it would *seem* that Protagoras could simply *shrug off* something like IC\*. (Just recall that here Protagoras would be able to accept the key point of IC\* and then reply that it doesn't follow from IC\* that PGR is *not true for Protagoras*.)

But having seen now that PGR amounts to not *merely* the 'no-mistakes' direction of the biconditional, but the *no-unbelieved-truths* direction as well, Burnyeat thinks we're (at least nearly, though not just yet) positioned to suggest that IC\* is under closer inspection a serious problem for the kind of global relativism Protagoras is peddling. The explanation is remarkable, albeit complex.

On Burnyeat's analysis, it's without a doubt the case that Protagoras is committed to saying that if one side of the PGR biconditional is true then the other is. But is Protagoras, in endorsing the biconditional, insisting that the predicates 'believes that *p*' and 'is true for *p*' *mean* the same thing?

This is puzzling. As Boghossian (2006b, 19) has remarked:

Relativism cannot properly be seen as correcting our view of what our sentences mean; it must rather be seen as correcting our view of *what the facts are*.

Burnyeat sides with Boghossian, as well as with Meiland & Krausz<sup>33</sup> (1982, 4) on this point and against Passmore (1970) and Grote (1865) insofar as he denies that the Protagorean relativist's commitment to the biconditional captured by PGR comes as well with a commitment to the

'synonymy thesis' that to say that 'x is true for S' is really just to say that 'S believes that x'. What Burnyeat finds objectionable is the thought that 'Protagoras [is] saying no more than that in discussing any proposition, the Measure doctrine included, all anyone can do is *express his own conviction*'.<sup>34</sup> Here Burnyeat's remarks are illuminating:

Protagoras's theory is, after all, a *theory of truth* and a theory of truth must link judgments to something else – the world, as philosophers often put it, though *for a relativist, the world has to be relativized to each individual*. To speak of how things appear to someone is to describe his state of mind, but to say that things are for him as they appear is to point beyond his state of mind to the way things actually are, not indeed in the world *tout court* (for Protagoras there is no such thing), but in the world as it is for him, in his world. (*Ibid.*, 181, my italics)

This passage is striking. And for the purposes of understanding the kind of beast PGR really is, it's important. For if the equivalence expressed by PGR is to be understood as taken by the Protagorean global relativist as a claim of mere *synonymy*, then it looks as though we make the mistake (as Burnyeat notes other commentators have made) of attributing to Protagoras no more than a *tautological* thesis, on which to say of one that one judges that *p*, one is saying nothing *more*, and *a fortiori* nothing more substantive, than *that* one judges that *p*. If this were right, then by Protagoras's lights, Protagoras would mean the same thing as Socrates does when the former says 'X is true for S' and the latter makes the unremarkable observation that 'S believes X'.

Where are we then if we agree with Burnyeat and deny that a PGR proponent is just trying to just tell us that 'X is true for S' means 'S believes/judges that X'? In the above passage, Burnyeat is expressing the idea that to say of Socrates that things *are for him* as he believes is to point to 'the way things actually are, in the world as it is for him, *in his world*'.<sup>35</sup> And notice that this is a kind of *metaphysical* claim, and a curious one at that. What would such a metaphysical claim involve? Alternatively: suppose we wanted to take it seriously. How are we supposed to do that?

Though Protagoras never explicitly said, Socrates imagines on his behalf, supposing that, as such a story must surely be part of the view he's advancing, he *must* have such a line, even if he perhaps only told his 'inner circle'. The unpacking of this metaphysical element of Protagorean global relativism has been termed Plato's 'secret doctrine', one which as Lee (2005, 77) remarks, represents Plato's attempt to build

as strong a case as possible for Protagoras by drawing upon additional resources which Protagoras himself did not use.

The 'secret doctrine' – viz., the 'final piece' of the Protagorean relativist's thesis – a piece needed in order to appreciate the full idea captured in large part by the biconditional – attempts to make sense of what this relativistic world would be like, one on which we are to suppose the PGR biconditional holds. And in order to make sense of such a world, Plato makes use of a perplexing Heraclitan analogy – one meant to illuminate 'an ontological setting' where each of us lives in a *private* world.<sup>36</sup> Such private worlds are constituted by a succession of 'momentary appearances' and of these momentary appearances, as Burnyeat remarks:

all...are true in that world quite independently of what happens next in a given world. In a given world – say, that of Socrates – whatever appears to him is then and there the case [sic. *no-mistakes*] and nothing is the case unless it then and there appears to him [sic. *no-unbelieved-truths*]. Such is the kind of world presupposed by Protagoras' doctrine that each man is the measure of all things. (Burnyeat 1976, 181–2)

This is, of course, a mindboggling idea. But this much seems clear: Plato thought it would be even *more* puzzling (than hypothesising such a world) to try to make sense of the thesis submitted on behalf of the global relativist *unless* we attribute to her something like the striking kind of anti-realism that by way of illustration the secret doctrine invites us to imagine.

A quick clarification here might be helpful, between the kind of ontological setting being described here, and the kind of commitment one incurs when one adverts to a 'possible world' where certain things are claimed to hold. The clarification is needed because in the latter kind of case, world-talk needn't carry with it any 'exotic ontology'. For example, as Theodore Sider (2002, 279) observes about the use of Kripke models to establish formal results in modal logic:

These models contain sets often regarded for heuristic reasons as sets of "possible worlds". But the "worlds" in these sets can be anything at all; they can be numbers, or people, or fish. The set of worlds, together with the accessibility relation and the rest of the model, is used as a purely formal structure.

By contrast, Burnyeat's suggestion is that taking PGR seriously involves attributing to Protagoras the thought that, if (for instance) numbers, or

people, or fish are believed by Socrates to be present, then *this is the case in Socrates' world* – viz., that Socrates' world is a number-world, a people-world and a fish-world.<sup>37</sup> And that it's Socrates' judging that *grounds* his world such a way. This is much more provocative than the thought that 'X is true for Socrates' means that Socrates judges that X. Though an interpretation of the thesis on which it comes out *as* provocative in its implications seems to align with what we should expect: global relativism is supposed to be a provocative thesis after all.<sup>38</sup>

At any rate, on the Burnyeat line, the attribution of such an ontological setting for each individual (where truth answers to belief, viz., where *man is the measure*) is needed to unpack the full expression of the Protagorean global relativist's thesis, else we're stuck relegating the discovery Protagoras purports to have made as just one of synonymy. Note, to be clear, that the ontological setting imagined by Plato in his recourse to the Hericlitian analogy is (obviously) just an illustrative example of how the anti-realist commitment of global relativism might be articulated. The point Burnyeat is after is just that any plausible way to interpret what the global relativist is putting forward will require *some* kind of profound anti-realism, of which the Hericlitean metaphor serves to illustrate but one example.

One might be inclined to (perhaps recoiling from the headspinning ontological setting intimated by the secret doctrine) insist that Protagorean global relativism must somehow be fully articulable, but without 'private-word anti-realism', where private-world anti-realism picks out the idea that each individual occupies a private world. But retreating in this way looks like a dead end.

Indeed, it looks like the attribution of this very blatant departure from what we saw in Chapter 1 as the independence leg of realism is a necessary ingredient of PGR unless we either:

- (i) articulate PGR in terms of only (NM), dropping (NUT) from the equation;<sup>39</sup> or
- (ii) articulate PGR in terms of both (NM), and (NUT) but take the Passmore–Grote line and attribute to the Protagorean global relativist a thesis that does not outstrip a claim of synonymy.

Option (i) is not good. One helpful way to see why is *via* a case Plato appeals to involving two individuals in the wind, an example Burnyeat (1976, 178) cites as compelling evidence that Plato took NUT (and not merely NM) to be essential to PGR. In the example, we are to imagine two people standing in the wind, where only one of the two individuals

feels cold. Plato very plausibly understands Protagoras's line on such a scenario as one on which the wind is 'cold for one of them... not so for the other' (p. 187, cf. 152b). Obviously, from NM, it follows that the wind is cold for the one who feels cold, though crucially the converse, NUT, is needed to get the result that 'it is not cold for the second person' (178). As Burnyeat analyses this point:

for it was not said that the second person feels warm in the wind, or even that he feels the wind is not cold; given either of these as premises, to conclude the wind is not cold for him we would need no more than...*If it seems to x that not-p, then it is true for x that not-p* which is a simple substitution of [NM]. What was said of the second person was simply that he *does not feel cold* (my italics). On Protagoras's view, then, if the wind does not appear cold to someone, that is sufficient grounds to assert that it is not cold for him, and this means that we must include [NUT]...in any complete formulation of the doctrine that man is the measure of all things'. (*Ibid.*, pp. 178–179)

Option (i) thus reduces significantly – and *implausibly* so – the scope of what the global relativist should be regarded as imputing in individual cases, by denying that the thesis entails such verdicts as 'X is not F for S' when S does not regard X as F.<sup>40</sup>

Option (ii), like option (i) avoids the strong (private-world) anti-realist reading of PGR and without dropping NUT from the equation, but it's not at all compelling why option (ii) should be thought of as offering the *lesser* of the two evils. For some *prima facie* provocative thesis T, a reading of T on which it comes out as (very) provocative is better than an interpretation on which this apparent provocativeness is explained away.

### 2.3.4 Protagorean global relativism and self-refutation reconsidered

Armed now with a suitable appreciation of the role NUT is playing in PGR, and the kind of strong anti-realist picture that looks needed to make sense of the full expression of PGR, it's time to make good on the promissory note, as it will be comprehensible now why Burnyeat took Plato to have been *right* in drawing attention, as he did, to the fact that an opponent of Protagorean relativism (i.e. Socrates) disbelieves PGR, and then to reason from here to the conclusion that Protagorean global relativism can, *on its own terms*, be turned against itself.

Let us pick up now with the formulation of the interim conclusion that (as was shown) Protagoras *must* accept, in light of his own doctrine and in light of the dialectical rule of acknowledging one's opponent's statement:

IC\*: *It is not true for Socrates that: if Socrates believes that  $p$ , then it is true for Socrates that  $p$ .*

Now, *originally*, it seemed that Protagoras could simply embrace the interim conclusion he is forced in to and retort that the doctrine remains true *for him* (and that Socrates' having thought otherwise must be down to Socrates failing to appreciate that Protagoras did not endorse an unrelativised truth predicate). That was at any rate the superficial lay of the land. But an apt appreciation of what IC\* involves, Burnyeat (1976, 186) thinks, closes off this escape route – by revealing how it is that Protagoras's adverting to his own belief in his theory 'counts for naught'.

We've already seen how Burnyeat's analysis exposed the anti-realist element of PGR as an indispensable element of the view. And soon we'll be in a position to see how, in construing IC\* as, despite initial appearances, something ultimately damning to Protagoras, a point of contemporary contention between Wright and Boghossian surfaces and which bears crucially on the matter of how PGR should be thought of as self-refuting.

In embracing IC\* Protagoras is *by his own lights* committed to accepting that in Socrates' world it is not a sufficient condition for the truth of a proposition that Socrates believes that proposition. And as Burnyeat observes, because Socrates also does not believe that 'If he (Socrates) doesn't believe something, then it's not true for him', Protagoras must grant that, in Socrates' world, neither is it a necessary condition for the truth of a proposition that Socrates believe that proposition. Generalising from the case of Socrates, it's clear that *if* no one believed the Measure doctrine, then, by the Measure Doctrine 'no one lives in a world in which his mere belief in a proposition is either a sufficient or necessary condition for its truth (in that world)' (*ibid.*, 182). And this is in fact, the first of three arguments Socrates makes (*viz.*, by insisting that if no one believes Protagoras's theory, then it would be true *for no one*) against PGR, as part of a 'triple sequence' of arguments that concludes with the *peritrope*.<sup>41</sup> More carefully: if no one believes Protagoras's doctrine, then by Protagoras's own lights, no one lives in the kind of world that the Measure doctrine asserts that *everyone lives in*.

But – and crucially – even if just *someone* believes the measure doctrine to be false, then, as Burnyeat puts it:

on Protagoras' own showing such persons do not, *as Protagoras alleges we all do*, live in a world in which their mere belief in a proposition is a sufficient and necessary condition for its truth (in that world) ... he is committed, despite himself, to agreeing that [his doctrine] is false.<sup>42</sup>

In sum, Burnyeat's analysis of the *peritrope* is that it gets Protagoras to accept (IC\*) – viz., that his theory is false for others – *but admitting this amounts to admitting that not everyone is a Protagorean measure* (which itself involves a commitment to acknowledging that there are some for whom belief in a proposition is neither necessary nor sufficient for the truth of that proposition). But *this* conflicts with what Protagoras is selling in submitting PGR, which is that everyone lives in a world in which his mere belief in a proposition is either a sufficient or necessary condition for its truth (in that world). As he remarks:

Hence it follows from Protagoras' admission that his theory is false for others *that it is false for himself as well*. There is a passage from "*p* is false for Protagoras' opponents" to "*p* is false for Protagoras" – *in the one special case where p is the Measure doctrine itself*.<sup>43</sup>

Now, then: does the Burnyeat reading of the *peritrope* argument (replete with the anti-realist construal of PGR) render the *peritrope* effective in showing PGR to be dialectically self-refuting? It is, advantageously, an interpretation on which we don't ostensibly attribute to the Protagorean global relativist an unrelativised truth predicate in order to force self-refutation. As a matter of scholarship, this interpretation does involve attributing to Plato a bit of carelessness in presentation (carelessness that is, at least as Burnyeat sees it, much more plausible than attributing to him an wholesale lack of understanding of Protagoras's position); though again, this interpretative point isn't what's at issue for our purposes. What matters is whether the reinterpretation of the argument shows the global relativist's view to be self-refuting (dialectically) in a decisive way.

With this in mind, consider that, crucial to the thought that Protagoras's assertion of his doctrine (in conjunction with his dialectical commitment to accepting (IC\*)) leaves him contradicting himself is the thought he is in fact committed to accepting *and* denying the following – that

everyone lives in a world in which his mere belief in a proposition is a necessary and sufficient condition for its truth (in that world).

But is Protagoras committed to regarding *this* claim as merely relatively true, or by contrast, is Protagoras submitting his thesis as one he is purporting to hold for everyone? It is without question that Burnyeat thinks (and thinks that Plato thought) Protagoras meant the *latter*. In fact, Burnyeat remarks that, whereas the absolute prefix 'It is (absolutely) true that ...' can be iterated over and over, a relativistic prefix such as 'It is true for Protagoras that ...' admits of only limited reiteration. At some point, though we may not be able to say just where, Protagoras must stop and take a stand.' Failing to do so, he thinks, will be at the cost of 'losing grip' (p. 194) of relative truth.

Accordingly, Burnyeat submits that at the end of the day, PGR is a thesis that must be understood as maintained in conjunction with what he calls the 'principle of translation':

*Principle of Translation* (POT): A proposition of the form 'x is F' is true (relatively) for person (a), if and only if, 'x is F for a' is true (absolutely).

Now if PGR should be understood as tempered (to the end of maintaining sensibility) by (POT), one interesting implication is that Plato's style of refutation advanced in the *Theatetus*, (at least, on the Burnyeat interpretation)<sup>44</sup> looks *considerably* more damning against the global relativist than does the line of reasoning widely assumed to be a 'knock-down' argument against (any version of) the view.

Hales (1997) characterises, pithily, the 'knock-down' rationale<sup>45</sup> as follows:

the relativist thesis is that everything is relative (nothing is absolute). Well, what about the claim itself, that everything is relative? It must be relative too – relative to a perspective, conceptual scheme, viewpoint, or what have you. In other words, there are perspectives in which the relativism thesis is true, and there are those in which it is untrue. After all, its truth is relative. Hence there is a perspective in which absolutism is true. This seems like a paradox, or a contradiction, or something.<sup>46</sup>

PGR, tempered by (POT) seems to evade the knock-down. The reply is that: we should understand the proponent of global relativism as submitting a claim that, *via* the principle of translation, *depends* upon

the countenancing of absolute truths (in particular, absolute truths that take the form of 'X is F relative to S'). So it is therefore not an argument against the position to draw attention to the fact that the global relativist is committed to absolute truths, *per se*.<sup>47</sup>

Furthermore, the appeal to POT cuts off at the pass the more *general* strategy of philosophical argument against PGR of which the knock-down argument is an instance. The general strategy, which is the *self-referential* strategy, proceeds, as Hájek (2014, 297) puts it, as follows:

take a philosophical thesis, and to make it refer to itself, to plug into a function *itself* as its own argument, and more generally, to appeal to self-referential cases. This technique is another handy way of cutting down the search space when you are looking for counterexamples.

Obviously, this is precisely the move that is made when one begins this way against the global relativist: 'Everything is relative? What about *relativism*?'

## 2.4 Global relativism and the principle of translation

To this point, I've said Burnyeat's diagnosis *reveals* both (i) the importance of reading the Protagorean global relativist as a strong anti-realist, and further, (ii) *that* Burnyeat's attempt to save Plato's argument highlights how adverting to the translation principle seems indispensable – a principle the status of which I suggested features in a more contemporary dispute between Wright and Boghossian. We'll consider (briefly, and foreshadowing Chapter 6) what's at stake between Wright and Boghossian on this point shortly, as it obviously bears on the plausibility of the Burnyeat–Plato rationale for why PGR is dialectically self-refuting.

But first I want to throw a spanner in the Burnyeat line that has gone unnoticed – and one that (like the matter of whether the translation principle is correct) bears directly whether the argumentative strategy considered in the previous section would be regarded as successful. The worry I raise will also have some implications for the Wright/Boghossian debate that centres around the translation principle itself.

### 2.4.1 Irreconcilable worlds objection

We can call my spanner, for lack of a better name, the 'irreconcilable worlds' objection,<sup>48</sup> one which highlights a tension between two key steps in Burnyeat's attempt to interpret the *peritrope* as a successful self-refutation of PGR.

The argument for such a tension runs as follows. The strong anti-realist reading of the PGR biconditional appeared necessary to avoid, as I think Burnyeat was right to note, a synonymy reading of PGR; the only other alternative was to drop NUT from PGR and read PGR as *just* NM; though this effectively ‘amputated’ PGR, leaving no explanation for why it should be thought to entail the kinds of assessments (*a la* the wind analogy) that are intuitively regarded as part and parcel with what the global relativist is submitting. Thus, the strong anti-realist reading of the NUT element of PGR appears to be *needed* for a plausible interpretation of the full expression of the doctrine. And I think Burnyeat is in fact right about this.

However, Burnyeat’s appeal to the principle of translation (POT) looks to be in a kind of *fundamental* tension with strong anti-realist reading of PGR that he rightly regarded as necessary to the end of an apt interpretation the view. The tension is both *metaphysical* and *motivational*.

Here, in a nutshell, is the metaphysical tension. The strong anti-realist construal of PGR (inspired by the secret doctrine) gives the notion of *private worlds* centre stage in attributing to Protagoras the view that everyone lives in a world in which his mere belief in a proposition is a necessary and sufficient condition for its truth (in that world). But if, as Burnyeat is suggesting, we *also* take Protagorean relativism to be presupposing POT, according to which a proposition of the form ‘*x* is *F*’ is true (relatively) for person (*a*), if and only if, ‘*x* is *F* for *a*’ is true (absolutely), then (obviously) with respect to the absolute truths (that take the form ‘*x* is *F* for *a*’) the truth of these depends *not* on how things are in an individual’s private world (else they not be absolute truths) but rather on how things are *independent* of one’s private world. But take now, for example, a counterfactual conditional:  $\phi \square \rightarrow \neg \psi$ , where (for instance) ‘*X* is *F*’ features in the  $\phi$  place in the antecedent, and we plug ‘*X* is *F* for *S*’ in for  $\psi$  in the consequent. For some agent *S*, the truth of the antecedent depends on *S*’s situatedness in *S*’s private world; however, the truth of the consequent (a truth that by POT will be *absolute*) must depend on *S*’s *not* being situated in a private world.

In short, then, the tension between Burnyeat’s two key moves in saving Plato’s *peritrope* against PGR has as a consequence that there will be cases where a subjunctive conditional’s truth conditions demand an agent to, *impossibly*, be in a private world and not be in a private world. This example case is but one way to make the more general metaphysical point that allowing some things to be absolutely true for one whose *world* is regarded as private (in the strong anti-realist sense that seems needed to give a full articulation of PGR) carries incoherent metaphysical

consequences. So this is a *prima facie* reason to think that: if we are to take seriously the thought that strong anti-realism is needed to avoid the synonymy reading (and other associated worries to do with what PGR can be understood as implying), then we can't as easily as Burnyeat does help ourselves to the POT as a plausible way to characterise what the view is claiming.

The *motivational* problem, like the metaphysical problem, concerns the tension in the view Burnyeat is attributing to Protagoras, a tension between (i) strong anti-realism (in unpacking PGR) and (ii) POT. The crux of the motivational problem is this: the philosophical motivation for PGR would, on the private-world reading, be so provocative that it is very strange to imagine *how* it could be germane to accounting for *some* truths but not others. Any motivation for such a view would seem to be one that would extend to truth, *tout court*.

This is a very abstract consideration, but we can put it plainly by considering an example. Suppose Socrates says 'The grass is green'. With reference to Burnyeat's strong anti-realist unpacking of PGR, the Protagorean global relativist, in accounting for how this statement is true *for Socrates*, postulates a private world in which his mere belief in a proposition is a necessary and sufficient condition for its truth (in that world). So if 'The grass is green' is true for Socrates, then Socrates lives in a private world (one in which that proposition is true for him because he believes it.) However, it seems that once this claim is made about Socrates, there's no going back – no describing his ontological setting differently when it's time to (*a la* POT) talk about a different truth, the related 'The grass is green is true for Socrates' which by POT is on (Burnyeat's thinking about) PGR supposed to be *absolutely* true. As with the metaphysical tension, then, we reach the dilemma that either Socrates' world is a private world, or not. And for one who *motivates* an anti-realist case to answer the matter affirmatively, (as we've seen there is independent reason to do) it would seem at best *ad hoc* to suppose one can conveniently backtrack and allow Socrates to escape his private world in order to account for a range of *a la carte* truths, the ones that are, by reference to POT, the absolute ones corresponding to each relative truth (by the formula specified by the translation principle).

Burnyeat's attempt to save Plato's own refutation thus runs in to a jam: either give up the strong anti-realist reading of PGR at the cost of implausible consequences, or maintain this reading only by facing a very troubling tension when imputing to the global relativist a view that *also* countenances the principle of translation.

Of course, one way out of the jam Burnyeat has gotten himself in to is to simply abandon the principle of translation when interpreting the claim being put forward by the global relativist – *viz.* by not claiming that making sense of relative truth requires it. We've yet to consider just how plausible such a move would be.

#### 2.4.2 Contemporary debate: Boghossian versus Wright

*Question:* to make sense of relative truth, do we *need* the translation principle? Obviously, Burnyeat (1976, 192–93) says *yes*. Interestingly, so does Boghossian (2006a). And both take this point to be an obvious one. By contrast, Wright (2008b) says not only that the relativist doesn't *need* the translation principle to make sense of relative truth, but even more, the very thought that the relativist *does* need to resort to it is to simply not take the idea of *relative truth* seriously.<sup>49</sup>

Considering the stakes, the translation principle marks what is perhaps the most deep divide in contemporary thinking about relativism and not one that we can very quickly resolve, though the issue will re-emerge in more detail Chapter 6. Perhaps the most helpful way to think about this contemporary divide for now is in terms of two very different ways of filling out the details of *arity* (degrees of relationality) thinking about relativism itself.

On one version of the arity approach, which we can call *relationism*, explicit reference to some parameter is needed in order to express the sense in which the relativist about some property P can submit P-claims to be true. The *relationist* approach to arity thinking about relativism owes principally to Boghossian (2006b), who envisions the relativist submitting that since (for the relativist) no propositions of the form 'X is P' are *absolutely* true, 'The closest truths in the vicinity are related *relational truths* of the form "X is *p* relative to F" where "F" names some appropriate parameter'. And this is more or less a restatement of the translation principle.<sup>50</sup>

Whereas Boghossian's relationist model conceives of the relativist (about some property *p*) submitting that claims of the form 'X is *p*' are *elliptical* for longer claims which themselves aspire to (absolute) truth, the (for lack of a better term) *non-relationist* (e.g. Wright 2008b) denies that for 'X is *p*' to be relatively true, on the relativist's framework, this involves smuggling the hidden parameter (F) into the (explicitly relational) *content* of a claim that is itself then put forward as absolutely true.

Here's now a very different way for the arity thinker to theorise about the role the unexpected parameter might play *vis-à-vis* relative truth.

Compare *X is p relative to parameter F* with the claim that '*X is p*' can be true or false, 'albeit only *relatively so*', that is, only once some relevant parameter (e.g. a *standard*) is specified or supplied. This *non-relationist* idea, on which explicit reference to the parameter is not a part of the content that is, for the relativist, a candidate for (relative) truth, has been articulated by among others Wright (2008b) and MacFarlane (e.g. 2014). On the Wright–MacFarlane model, (what Wright calls 'new age relativism') the surprise the relativist highlights *isn't* that any rendering of '*X is p*' as true must involve treating *p* is a two, rather than, one-place *predicate* (e.g.  $\langle p, F \rangle$ ). Rather, the 'surprise' is that utterances of '*X is p*' take the specification of a *standard* to get a truth-value.

Compare here: the foundational work in semantics by Lewis (1980) and Kaplan (1989) have shown that we (already) need to relativise truth to world, time, location triples.<sup>51</sup> Though this isn't really surprising given the (obvious) non-specificity of contents and operators in a language. Rendering the truth of moral utterances as dependent upon the specification of a standard (as supplied by what MacFarlane calls a 'context of assessment') is by contrast much more unexpected. The semantic machinery deployed by the MacFarlane-style relativist is in fact much more complex than has been sketched so far. And we will be engaging with this version of relativism in much more detail in Chapters 7 and 8. For now, the relevance of drawing attention to the view is to show how there are various ways of unpacking the notion of an unexpected parameter, in characterising relativism in terms of arity; in short, one of the two ways embraces the principle of translation, the other does not.

## 2.5 Global relativism, whither now?

It looks, then, like the translation principle Burnyeat appeals to in attributing to the global relativist what he takes to be a comprehensible position is – along with, as I've suggested in the irreconcilable-worlds line of critique, being at odds with a plausible anti-realist reading of PGR – very much *theory-laden* in that it takes for granted a particular and contentious way of thinking about relativism, in terms of just how to unpack the arity insight – one that is very much a matter of debate between Boghossian and those such as MacFarlane and Wright whose interest in relativism comes more squarely from the semantic side of the debate.

I want to conclude by drawing attention to how, *regardless* of whose side we take in the Wright/Boghossian debate *vis-à-vis* the translation

principle – and more generally, where to locate the unexpected parameter – things do not end up looking good for the *global* relativist.

As we've already seen, Burnyeat makes a compelling case that *if* we attribute to, in the case of Protagoras's 'man-parameter' version of global relativism, the proponent of PGR both the strong anti-realist version of the thesis and then attribute to her, as well, a commitment to accepting the translation principle, then Protagoras is stuck dialectically refuting himself in asserting his doctrine, in virtue of, *a la* Burnyeat, incurring an endorsement and a denial of the claim that everyone is a Protagorean measure (e.g. such that everyone lives in a world in which his mere belief in a proposition is a sufficient and necessary condition for its truth (in that world)).

And, given the plausibility of reading PGR in the strong anti-realist way, things thus do not look good for the prospects of global relativism insofar as we follow Burnyeat and Boghossian and attribute to the global relativist the translation principle as underlying her thought (even though the irreconcilable-worlds objection revealed such an attribution an implausible *characterisation* of the view when *paired* with the anti-realist reading of PGR).

*Question:* does the global relativist come out faring any better if we do *not* attribute to her a tacit commitment to the translation principle? Let us suppose we do not. In this case, the truth that the Protagorean doctrine *itself* aspires to is relative truth. Crucially, we are not, on this line of thought, tacitly supposing that by a 'relative truth' we mean: 'Protagorean relativism is true relative to Protagoras' is absolutely true. Rather, we mean that the claim Protagoras expresses in avowing his theory is true relative to himself, though not relative to anyone (for instance, his opponent, Socrates) who does *not* believe it.

Taking this line of thought further, though, a hitch materialises, one that I think shares some commonalities with a related hitch that Wittgenstein (1969) highlights in *On Certainty*, when expressing his dissatisfaction with the thought that one might, *a la* Descartes, call in to doubt the *totality* of one's beliefs at once.<sup>52</sup> In pursuing this parallel, consider that, in what is now a famous passage, Wittgenstein (1969, §§341–3) remarks:

[...] the questions that we raise and our doubts depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.... That is to say, it belongs to the logic of our scientific investigations that certain things are indeed not doubted. But it isn't that the situation is like this: We just can't

investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put. (OC, §§341–3)

The propositions Wittgenstein famously declared *arational* in the sense that we can neither rationally support them with other more certain propositions, nor rationally doubt them, are termed ‘hinge propositions’.<sup>53</sup> Granted, the *status* of such propositions is very controversial among contemporary commentators on Wittgenstein’s epistemology, and we needn’t get caught up on this point. The general insight that will be relevant to us, put simply, is that the very practice of offering and giving reasons for belief requires that some things be taken for granted. The rationale for this insight is grounded in Wittgenstein’s account of the structure of rational support relations, and the key idea that will interest us here is, as Pritchard (2011, 194) puts it, that:

in order for something to be a ground for doubt it is essential that it be more certain than that which it is calling into doubt, since otherwise one would have a better epistemic basis for rejecting the ground for doubt than for rejecting the belief which is the target of the doubt.

Consider an example Wittgenstein offers in illustrating this point:

If a blind man were to ask me “Have you got two hands?” I should not make sure by looking. If I were to have any doubt of it, then I don’t know why I should trust my eyes. For why shouldn’t I test my eyes by looking to find out whether I see my two hands? What is to be tested by what? (OC, §125)

Let us suppose that Wittgenstein has cottoned on to something correct about the rationality of *doubt*. If so, then there is very plausibly a parallel kind of Wittgenstenian-style point that can be drawn out in the service of raising a problem for the global relativist who shuns the translation principle and thus submits that that the claim being put forward by the global relativist aspires to truth, *albeit merely relative truth*. To reiterate, ‘relative truth’ here is not being understood, *a la* Boghossian, as an absolute truth about a proposition of the form ‘S is F relative to parameter P’. We are rather working with the hypothesis that global relativism aspires to relative truth in the Wright/MacFarlane sense: that the content expressed when the global relativist submits her view gets a truth value

only relative to some parameter which is specified in a context where that content is assessed for truth or falsity.

Now it would seem very plausible that the very practice of *relativising* one thing to another requires that some things are not relativised, much as (as Wittgenstein suggests) the very practice wherein rational doubt is possible requires that some things are indeed *not* doubted. In the case of rational doubt, Wittgenstein's example at (OC, §125) is meant to illustrate the non-sensicality of thinking otherwise. I think we can tease out a parallel example to illustrate the kind of problem that crops up for the global relativist strategy currently under consideration.

Suppose the global relativist (of the variety under consideration) does *not* take for granted that there are *non-assessment-relative facts* about (to revisit the locution used by the co-variance approach) the *domain of relativisation*, and its relationship to the object of relativisation. Again, in the case of PGR, the domain is 'man' (as we've unpacked this), and the object of relativisation is *all truths*, including truths *about* what is being relativised to what.<sup>54</sup> Suppose now Socrates judges something to be false: he asserts 'X is false'. And the Protagorean global relativist says 'X is false' is true *for* Socrates because Socrates believes this. But what is the status of the claim operating in the background, that Socrates believes what he purports to believe e.g., that X is false?

The unanchored global relativist under consideration tells us that *whether Socrates believes that X is false is itself an assessment-sensitive matter*: 'Socrates believes that X is false' gets a truth value only relative to a context of assessment. Belief does not automatically iterate; it is very possible that Socrates just believes that 'X is false' but does not believe that he believes that X is false. But then, if he does not, then it's *not true for Socrates that Socrates believes that X is false*. But if *that's* right, then the unanchored global relativist's original story about Socrates' belief that 'X is false' unravels; that story recall was premised upon the suggestion that Socrates' belief that 'X is false' is *true for Socrates because Socrates believes that X is false*. But we've just established that by the unanchored global relativist's own lights, the causal explanation here (e.g. that Socrates believes that X is false) might be one that is *false for Socrates, even if Socrates believes that X is false*.

It is hard at this point not to feel the pull of the Wittgenstein-style point alluded to earlier; just as the process of rational doubt requires that some things are in *fact* not doubted, it looks very much like a coherent picture on which one thing is relativised to another, requires that *some things are not relativised*; at the very least, some things (in particular, facts about that which the object of relativisation is relativised *to*) must be taken

for granted in such a way that they are not regarded as of the same *degree of relationality* (arity) as the *object* of relativisation itself. Else we end up in incoherence, as in the kind of situation just described, which is no less nonsensical than the corresponding example, in the practice of doubt, where something is submitted as a grounds for doubt which is less certain than that which it is calling into doubt.

## 2.6 Concluding remarks

So where does that leave us, regarding the plausibility of the *trivial entailment* argument for epistemic relativism – a royal road from an independently established global relativism? From what we've seen here, there really is no plausible way to get to epistemic relativism *via* global relativism. But this is *not* because global relativism succumbs to some simple 'knock-down' argument as is, one gets the impression, a wide presumption. Compare: global *scepticism* is no more popular than global relativism, and global scepticism isn't a beast to be defeated by a simple knock down.

Rather, dismissing the global relativist takes some patience. In this chapter, several ways to think about what *relativism* is were canvassed, after which it was shown through a case study of Plato's attempted refutation of Protagoras just how a *dialectical* self-refutation of global relativism might proceed. Burnyeat's analysis of Plato's argument was revealing to this end, in that it teased out two important philosophical issues integral to an assessment of the plausibility of global relativism, and both of which will end up having some bearing on our subsequent discussions of (local) epistemic relativism in later chapters. Burnyeat's guiding insights, in particular, were that Plato's famous self-refutation argument against Protagoras' brand of global relativism succeeds (contrary to initial appearances) *provided* Protagorean global relativism is suitably appreciated as being committed to a strong form of anti-realism, which (as Burnyeat sees it) is then argued to be incompatible with dialectical commitments Protagoras incurs *via* what Burnyeat takes to be Protagoras's tacit commitment to the *translation principle* – a principle that (as we've seen) turns out to be front and centre in a contemporary debate between Boghossian and Wright on the matter of just how to think of relativism, of any stripe, *per se*.

While this was not the place to adjudicate between Boghossian and Wright on the plausibility of the translation principle (according to which 'x is F' is true relative to S iff 'X is F is true relative to S' is absolutely true), I argued that, *even if* an attribution of strong anti-realism

and the translation principle to the Protagorean global relativist would suffice (as Burnyeat thinks it does) to force Protagoras into dialectical self-refutation, a tension between these two elements ultimately renders Burnyeat's interpretation of the thesis Protagoras was submitting under the description of the measure doctrine unstable as a *characterisation* of any global relativist's position. This was the point of the irreconcilable-worlds argument.

Nonetheless, it was granted that *were* such a characterisation apt, then Burnyeat is right that dialectical self-refutation materialises for the proponent of Protagorean global relativism. And this means that the avenue left for any would-be Protagorean global relativist would be to *either* abandon a commitment to strong anti-realism (which I've argued would come at the cost of embracing either an implausibly 'amputated' version of PGR (in terms of just NM), or implausibly reducing PGR to a synonymy thesis) *or* to maintain strong anti-realism and abandon the *translation principle* – and in doing so embrace the thought that relativism *itself* aspires to merely relative truth, in the assessment-sensitive sense of relative truth intimated by MacFarlane and Wright. And as we saw, in drawing in the final section some parallels from Wittgenstein's *On Certainty*, this strategy, as well, was argued to be 'unanchored' in a way that turns out to be entirely unworkable. Global relativism is thus not a viable position, not because it can easily be shown to be false, but because there does not seem to be a tractable way to defend the thesis that does not ultimately lead the proponent to either dialectical self-refutation or unanchored incoherence. Obviously, from the unworkability of global relativism, it remains open whether *local* varieties of relativism are workable, and in the next three chapters, we'll consider the case for *epistemic* relativism on its own terms.

# 3

## The Pyrrhonian Argument for Epistemic Relativism

*Abstract.* This chapter critically evaluates one of the most provocative contemporary rationales for epistemic relativism, one which takes as a starting point the ancient Pyrrhonian problematic (Sankey 2010; 2011; 2013). It is shown though that an attempt to motivate epistemic relativism in this fashion falls flat despite its initial promise.

### 3.1 The Pyrrhonian argument for epistemic relativism

Epistemologists typically associate ‘Pyrrhonism’ with two key ideas:

*Pyrrhonian Scepticism* – viz., of the sort embraced by Sextus Empiricus in *Outlines of Pyrrhonism*; and

The *Pyrrhonian Problematic* – viz., a puzzle about epistemic justification, with reference to which foundationalist, coherentist and infinitist approaches to the structure of justified belief are characteristically motivated.

Pyrrhonian sceptics were, in fact, *not* relativists (epistemic or otherwise) and would regard anyone spouting a relativist doctrine as among the very kinds of *dogmatists* the Pyrrhonist movement was itself a reaction to;<sup>1</sup> the actual, historical Pyrrhonists<sup>2</sup> did not, as Michael Frede (1998a, 2) puts it, ‘purport to have any deep insights’, and relativism is a deep insight.<sup>3</sup>

Nonetheless, one of the most provocative recent arguments for epistemic relativism takes as a starting point the very same puzzle the Pyrrhonian sceptics wielded against the dogmatists, and with reference

to this puzzle, the thesis that epistemic relativism is the correct theory of epistemic justification is motivated. Howard Sankey,<sup>4</sup> the primary proponent of the view that the Pyrrhonian puzzle motivates epistemic relativism, goes so far as to say in his most recent paper that the core argument found in the Pyrrhonian problematic ‘constitutes the *foundation* for contemporary epistemic relativism’<sup>5</sup> (Sankey 2012, 184). It goes without saying that if Sankey is right, then the plausibility of epistemic relativism turns crucially on the plausibility of the ‘quasi-Pyrrhonian’ argument.

A quick note of clarification is in order before proceeding. Sankey himself is *not* a relativist. Sankey, after proposing the Pyrrhonian argument for relativism as what he takes to be the central way to capture the thrust of the relativist’s argument, attempts to show (e.g. in Sankey 2010) through a naturalistically inspired argument of his own how the relativist conclusion can ultimately be blocked. Because the most interesting aspect of Sankey’s work – and the theme that runs through his three papers – is that epistemic relativism *can* be motivated *via* the Pyrrhonian problematic (and indeed that this is, by his judgment the *best* way to motivate epistemic relativism), I am for ease of presentation writing *as though* Sankey *himself* is a relativist, though I reiterate that this is *just for ease of presentation*.

### 3.1.1 Pyrrhonian argument template

Suppose you claim to know that ‘*p*’, but your friend says ‘not-*p*’, and challenges you to defend your claim, asking you for a *good reason* for *p*. On the assumption that good reasons – e.g. the sort of reasons good enough to epistemically justify a belief – are *non-arbitrary* reasons, reasons that we have *good reason* to believe, a regress threatens: it quickly looks as though knowledge as well as epistemic justification require an *infinite* number of good reasons,<sup>6</sup> something it seems we don’t have, and so it looks like we don’t know or justifiably believe anything. Obviously, most epistemologists aren’t sceptics. And so something has to be rejected. But what?

Let’s unpack the explicit premises of this argument<sup>7</sup> – the ‘regress formulation’ of the Pyrrhonian problematic, as follows:

*Pyrrhonian Problematic* (regress formulation):

1. In order to be justified in believing something, one must believe it on the basis of good reasons.
2. Good reasons must themselves be justified beliefs.

3. Therefore, in order to be justified in believing something, one must believe it on the basis of an infinite number of good reasons.<sup>8</sup>
4. No human being can have an infinite number of good reasons.
5. Therefore, it is humanly impossible to have justified beliefs.

Premise (5) is a sceptical conclusion, and in fact a *stronger* sceptical conclusion<sup>9</sup> than the Cartesian sceptical conclusion, framed in terms of *knowledge*. So what gives? Everyone, except *epistemological infinitists*,<sup>10</sup> deny (3) of the argument. But since it looks like (3) is unavoidable if you accept (1) and (2), naturally, most philosophers deny either (1), (2), or deny that (1) and (2) entail (3) (see Lammenranta 2008 for discussion).

Historically, the most popular way out of the puzzle has been to simply deny (1) and insist that there are some beliefs – i.e. *basic beliefs*, which are justified but not *in virtue* of some relationship to another justified belief. This is the move made by *foundationalists*, a move embraced by Aristotle, the Stoics, Descartes, Hume and many contemporary externalists and internalists.<sup>11</sup> Coherentists, on the other hand, are divided. *Holistic* coherentists<sup>12</sup> deny (1) and say that a belief's status as justified is in virtue of its *membership* in a coherent set of beliefs rather than in virtue of any inferential relationship to some other particular belief. *Linear* coherentists by contrast deny that (1) and (2) entail (3) and accordingly allow justified beliefs to be supported by circular reasoning chains.<sup>13</sup>

One other prominent (albeit idiosyncratic) way to escape the sceptical conclusion takes inspiration from Wittgenstein's response to Moore (1939) in *On Certainty* – perhaps chains of good reasons can, at bottom, end in beliefs that themselves are 'arational', lacking the status of epistemic justification – e.g. what Wittgenstein called hinge propositions.<sup>14</sup>

Finally, the *sceptical infinitist*<sup>15</sup> responds to the puzzle by accepting its conclusion; the sceptical infinitist embraces the infinitist premise (3) and (unlike Klein (2007; 1999; 2003), who avoids scepticism by denying (4)) simply grants (4) and accepts (5)<sup>16</sup>.

### 3.1.2 Sankey's relativist redeployment

Notice that the Pyrrhonian Problematic surfaces once justification is *requested*. One place – familiar in epistemological theory – where justification is famously requested arises when one proposes a *criterion* for knowledge.<sup>17</sup> The 'Problem of the Criterion' looms: in order to recognise cases of knowledge, a *criterion* is necessary. But in order to determine whether a given criterion is a criterion of *knowledge*, there must be a way to recognise cases of knowledge *independently* of the criterion. (After all,

if we can't *independently* recognise cases of knowledge, we won't be able to tell that a given criterion is a criterion of *knowledge*, as opposed to something else). But where then to begin?

At this point, Sankey argues, we are led right in to the notorious Agrippa's Trilemma.<sup>18</sup> After all, it looks like our options are familiar ones:

*Option 1:* Adopt a criterion arbitrarily (without providing any reason for doing so).

But, as Sankey suggests, an obvious objection to Option 1 is that the adoption of the criterion in this manner would be *unjustified*.<sup>19</sup> Thus, it seems *prima facie* more promising to attempt to provide a justification for the criterion. But how is this to be done? How about:

*Option 2:* Justify the criterion by appealing to a further criterion.

However, appealing to a further criterion then raises the question of how the further criterion is to be justified. Iterate the process of appealing to further criteria, and what results is an infinite regress. Of course, there is another move:

*Option 3:* Justify the criterion by appealing to the criterion itself.

But, Sankey says, avoiding the regress this way is just to 'argue in a circle'.

The upshot, as Sankey (2010, 5) puts it, is that 'the attempt to justify the criterion leads either to infinite regress, circularity or unjustified adoption of the criterion.' The Pyrrhonian sceptic, in the face of these options, simply withdraws from attributing knowledge (and *mutatis mutandis* justification) to anyone; after all, in order to recognise cases of knowledge/justification, a criterion is necessary.

We are now at a familiar place: the argument Sankey's just run is, of course, just an *instance* of the Pyrrhonian regress argument considered in the previous section. It's at this point that Sankey thinks things can be taken in a *relativist* direction.

Firstly, to be clear, Sankey defines *epistemic relativism* as a view about *epistemic norms*.<sup>20</sup> He defines an *epistemic norm* as 'a criterion or rule that may be employed to justify a belief':<sup>21</sup> for example, the rule that it is (epistemically) acceptable to believe the deliverances of the sense perception. Epistemic relativism is, according to Sankey, the thesis that

there are no epistemic norms over and above the variable epistemic norms operative in different (local) cultural settings or contexts, where these local contexts are defined as always including at least a *system of beliefs*<sup>22</sup> and a *set of norms* (Sankey 2012, 187). Examples include not just local cultural contexts, but also system-of-belief/norm pairs featuring in Popperian frameworks, Kuhnian paradigms, etc.<sup>23</sup> For Sankey's relativist, whether a belief is justified, or counts as knowledge, depends on *epistemic norms*, and so, given that different epistemic norms can operate in different contexts, the same belief might be rational/justified/knowledge relative to one context, and not to another.

So, how does an argument *for* epistemic relativism, so conceived, materialise *via* tracing the steps that led the Pyrrhonian to scepticism? Let's now reframe things slightly, specifically in terms of epistemic norms, as Sankey defines them. Take an epistemic norm,  $N_1$ . Question: how is  $N_1$  to be justified? *A la* the Pyrrhonian puzzle, the options don't look very promising. One option is to Justify  $N_1$  by appealing to a further epistemic norm,  $N_2$ . Another option is to justify  $N_1$  by appealing to  $N_1$ .

Sankey says neither of these options satisfactorily vindicates  $N_1$  as justified; the former generates an infinite regress, the latter is viciously circular. Now: take *any other epistemic norms*,  $N_3, N_4 \dots N_n$ . By running through this *same line of thinking* with any of  $N_3, N_4 \dots N_n$  in an attempt to justify any of these norms, we end up in the same place. That is: each of  $N_1$  and  $N_3, N_4 \dots N_n$  are *equally lacking in justification*.<sup>24</sup> As Sankey writes:

If no norm is better justified than any other, all norms have *equal standing*. Since it is not possible to provide an ultimate grounding for any set of norms, the only possible form of justification is justification on the basis of a set of *operative* norms. Thus, the norms operative within a particular context provide justification for beliefs formed within that context. Those who occupy a different context in which different norms are operative are justified by the norms which apply in that context ... the relativist is now in a position to claim that epistemic justification is relative to locally operative norms.<sup>25</sup>

Sankey accordingly reveals that he thinks the Pyrrhonian problematic, *regardless* of what it might show in the service of promoting the Pyrrhonian verdict, suffices to show that it is not possible to provide a justification for choosing one epistemic norm over another.<sup>26</sup> And what this seems to imply is that all epistemic norms are *equally justified* – one's choice may rest upon such things as 'an irrational leap of faith, a subjective

personal commitment' or simply the convention of accepting the norms that prevail in whatever context one finds oneself.<sup>27</sup>

### 3.2 Two lines of objection

Sankey's idea can be helpfully broken in to two 'stages'. The first stage can be thought of as a kind of function/input/output model. If we think of the Pyrrhonian Regress as a function, which takes epistemic norms as arguments, the 'Pyrrhonian Regress function' maps norms to norms, where each output is the norm input as the argument, but *unjustified*.

As such, the Pyrrhonian Regress, *qua* function, might be called a 'justification eater.' No epistemic norm can be run through the regress with (any level of) justification intact. And since one unjustified norm is as justified as another unjustified norm, all norms (run through the function, at least) end up with the (relativist-friendly) epistemic status of *equally (un)justified*. See Figure 3.1.

The first stage of the argument, thus reaches the interim conclusion that all norms are of equal epistemic standing. This conclusion, in and of itself, does not imply epistemic relativism. The second stage of the argument thus must get us from the interim conclusion to epistemic relativism. On Sankey's own terms, this will be the conclusion that that there are no epistemic norms over and above the variable epistemic norms operative in different (local) cultural settings or contexts.

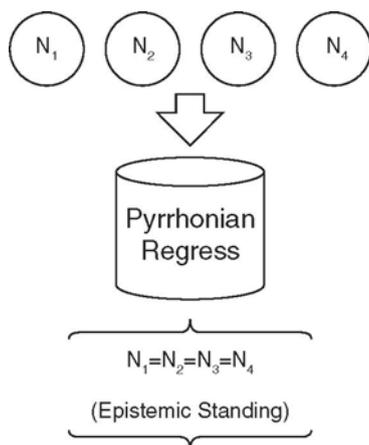


Figure 3.1 Sankey's argument: stage 1

In what follows, I offer two lines of objection, the first directed to the first stage of the argument – *viz.*, the argument that the Pyrrhonian regress motivates Sankey's intermediate conclusion that all norms are of equal epistemic standing. Next, I'll challenge the second stage of the argument and insist that even if we grant Sankey's (relativist's) intermediate conclusion, Sankey's given us no good reason to accept the conclusion that epistemic relativism is true.<sup>28</sup>

### 3.2.1 First objection

The careful reader will notice that when Sankey's relativist runs epistemic norms through the regress, in order to show that epistemic norms can't themselves be vindicated as epistemically justified, he's really only making explicit the unsatisfactoriness of two 'modes' of Agrippa's trilemma: *infinite regress* and *circularity*.

The way things are set up, escaping through the traditional foundationalist avenue (one which denies that justification comes only on the basis of good reasons) is not considered as a serious option – *viz.*, the option that some epistemic norm might attain the status of justification though not on the basis of any connection to some other epistemic norm.

This is odd. After all, foundationalists insist that something *x* can be epistemically *justified* even if *not* on the basis of some further thing *y* that one might cite as a reason for *x*. There are broadly two ways this might go. *Traditional* foundationalists allow that *x* might be *self-justifying* for the following reason: that the *truth* of *x* suffices to justify *x*. By contrast, and as Turri and Klein (2014, 6–7) note, *meta-justificatory* foundationalists hold a more sophisticated view: meta-justificatory foundationalists deny that it's simply the *truth* of *x* that does the relevant justificatory work; rather, *x* must have some 'further property *F*' (e.g. reliability). And metajustificatory foundationalists disagree about what the property *F* should be.<sup>29</sup>

It's hard to see how Sankey, in light of how he sets things up, is not begging the question against the *traditional* foundationalist. (Maybe that's not so bad – perhaps meta-justificatory foundationalism is the more plausible version.) But Sankey begs the question against the meta-justificatory foundationalist as well. As Fantl (2003, 541) has noted, whatever the 'F' is that on the metajustificatory proposal is supposed to be what gives *x* its status as foundationally justified, the metajustificatory foundationalist *can't*, as Fantl puts it, 'require that a believer have *access* to the metajustificatory feature as a reason for the foundational

belief.<sup>30</sup> As Turri and Klein (2014, 6–7) note, were the metajustificatory foundationalist to advert to such a requirement, this would:

undermine its putative status as *foundational*...It would effectively require a further reason for that which supposedly stood in no need of it.<sup>31</sup>

But once this point is made explicit, then it looks like, *in virtue of being a foundationalist view*, a view fails to count as a candidate to vindicate an epistemic norm as justified. If any version of foundationalism were true, Stage 1 of Sankey's argument faces a hitch: the Pyrrhonian Regress, taking a given epistemic norm as an argument, *doesn't* (so obviously) spit the same norm out without justification. It would do so *only if* there is no plausible (traditional or metajustificatory) foundationalist vindication of epistemic norms available. This is something that must be shown not just assumed.

Interestingly, the *holistic coherentist* is, like the foundationalist, in a position to object to the set-up of Stage 1. The holistic coherentist rejects that (put generally) in order to be justified in believing something (including some given epistemic norm), one must believe it on the basis of good reasons, of the sort one might supply to a sceptic who requests justification. (That is, holistic coherentists reject (1) of the Pyrrhonian regress argument.)

Let's bracket the points about foundationalism and holistic coherentism. Even if these approaches to vindicating epistemic norms as justified are incompatible with the presupposition of Sankey's set-up – that bona fide justification for a given claim *x* must come in the form of a further citable reason for the claim – it remains the case that *linear coherentists* as well as *infinetists* both embrace this presupposition behind the sceptic's request.

Let's assume for now that Sankey is right that linear coherentism is viciously circular.<sup>32</sup> There remains a problem with Sankey's dismissal of infinitism. The problem in short is that if infinitism is an untenable way to vindicate an epistemic norm as justified, it's not untenable for the reason Sankey most likely is assuming.

For in each of Sankey's three papers (2010; 2011; 2013), the infinitist thesis is regarded (implicitly) to be a *reductio against itself*, as opposed to a view that must be dismissed on the basis of some further unacceptable result that is not a component of the thesis but a consequence of embracing it. At any rate, this is the thinking we find in Sankey (2010, 5) and in Sankey (2012, 187), when what is highlighted is just *that* the infinitist approach leads to infinite regress. In Sankey (2011, 564), what

is said is that 'if one appeals to some further criterion, then the way is open to an infinite regress, since the further criterion must surely be justified, and so on *ad infinitum*.' But remember: the infinitist's *thesis* is that in order to be justified in believing something, one must believe it on the basis of an infinite number of good reasons (this was, recall, premise (3) in the Pyrrhonian Regress Argument).<sup>33</sup>

This would be an obviously invalid argument:

(I): *Infinitism*: In order to be justified in believing something, one must believe it on the basis of an infinite number of good reasons.

(~I): Therefore, (I) is implausible.

The best interpretation of Sankey's reason for rejecting that an epistemic norm could be vindicated as justified *via* infinitism is probably a tacit commitment to the *finite mind* objection – the most common – and yet, misguided – objection to infinitism.

*Finite Mind Objection*:<sup>34</sup>

6. We have finite lives and finite minds (Premise).
7. If we have finite lives and finite minds, we cannot produce an infinite series of reasons (Premise).
8. We cannot produce an infinite series reasons (From 1–2).
9. If infinitism is true, then in order to be justified in believing something, we must be able to produce an infinite series of reasons (Premise).
10. We have some justified beliefs.
11. Therefore, infinitism is not true.

If one takes infinitism to be a *reductio against itself*, it's most likely because one thinks that we simply can't have an infinite number of reasons, something that seems implied by the infinite regress. But infinitists do not actually embrace (9). As Turri and Klein (2014, 13) remark:

Rather, they typically say that we must have an appropriately structured, infinite set of reasons *available* to us... just having the reasons available, and producing enough of them to satisfy contextual demands, suffices to justify your belief.

And, as they note further, contextual demands rarely require that we be able to cite more than, for instance, ten reasons.<sup>35</sup> Even if contextual demands require that we be able to cite (say) 30 reasons (as, for example,

one might in a hostile interrogation), there would be no reason to write off infinitism on the basis of the most *obvious* objection to infinitism (one which reasons through premise (9)<sup>36</sup>). But, then, if the most obvious objection associated with an attribution of infinite reasons isn't enough to lay waste to infinitism, then it's not good enough to simply dismiss the view *as one that leads to infinite regress*. Infinitism can't simply be *assumed* to be unable to provide a vindication of the justification of epistemic norms.

Stage 1 of the argument is starting to look very shaky. In order to – by reference to the Pyrrhonian problematic – reach the intermediate conclusion that all epistemic norms are on equal standing (that is, equally *unjustified*) we need some non-arbitrary reason to think that *neither* the foundationalist, *nor* the holistic coherentist *nor* the infinitist could successfully vindicate an epistemic norm as justified.

### 3.2.2 Second objection

Let's revisit what Sankey says about the foundationalist-style strategy for defending an epistemic norm. Now I've already suggested that the relativist's rationale that Sankey presents does not take seriously the thought that an epistemic norm can have the status of being justified even if not in virtue of being accepted on the basis of good reasons.

Let's focus now on what specifically Sankey thinks is *wrong* with accepting an epistemic norm but *without* the provision of some further good reason one can cite. Here is a selection of three such remarks, found in three separate papers on the topic, which he's made on this point. Firstly:

If the regress is halted by the adoption of a criterion without justification, the criterion fails to be adopted on a rational basis.<sup>37</sup>

Here's another:

Alternatively, one might simply adopt the original rule dogmatically, without justification. But if the original rule is adopted in this way, it is adopted *without any basis*, and so is unjustified.<sup>38</sup>

And finally, another:

One way to respond to the regress of justification is to terminate the regress at a dogmatic halting-point. This may be done by simply adopting the criterion as an assumption... [if] a criterion is adopted by assumption, *it fails to be justified*.<sup>39</sup>

In each case, Sankey makes it clear *why* he doesn't take a foundationalist vindication of epistemic norms seriously – namely, he seems already convinced that Premise (1) in the Pyrrhonian Regress argument is *obviously* true.<sup>40</sup> That is, he thinks that (in the instance of (1) where the belief in question is a belief that some epistemic norm is true) in order to be justified in believing that an epistemic norm is true, one *must* believe that epistemic norm on the basis of good reasons. At any rate, being wedded to (1) of the regress argument would explain neatly *why* he makes the three remarks above.

The problem is, though, that if a commitment to (1) is supposed to be the underlying rationale for simply dismissing the foundationalist vindication of an epistemic norm (in Stage 1 of the argument), Sankey won't be entitled to then *violate* (1) when advancing the *second* stage of the argument, which moves from the intermediate conclusion that all epistemic norms are on equal standing to the conclusion that epistemic relativism is true. But I want to suggest now that that's precisely what he does.

To make this point, consider firstly, as Markus Seidel (2013, 137) has noted, that Sankey's envisioned relativist actually 'goes a long way with the sceptic'. Sankey will not dispute that his relativist travels the same road as the sceptic does. But Seidel is right to observe that Sankey travels with the sceptic long enough that he 'is at pains to provide us with reasons [for the relativist to] part company'. Once it's been claimed that all norms are equally *unjustified* – no norm is more justified than any other *in any way* – it's not apparent, as Seidel observes, how locally credible epistemic norms are supposed to have *any* positive epistemic status, positive status the relativist wants to *preserve* when insisting that epistemic norms aspire to relative justification.<sup>41</sup>

In short, even if we *accept* Sankey's intermediate conclusion that all epistemic norms are *equally unjustified*, it looks like the dialectical position favours relativism no *more* than it favours a move in the sceptical direction – *viz.*, a move from the intermediate conclusion that all norms are equally unjustified to the *withholding of judgment* about whether a given epistemic norm is correct.

Sankey's relativist's rationale is that having reached the intermediate conclusion that all epistemic norms are *equally unjustified* it follows that 'the only *possible* form of justification is justification on the basis of a set of *operative norms*' (Sankey (2012, 187)). Thus, he reasons:

the norms operative within a particular context provide justification for beliefs formed within that context. Those who occupy a different

context in which different norms are operative are justified by the norms which apply in that context...the relativist is now in a position to claim that epistemic justification is relative to locally operative norms.

This reasoning makes some dubious moves. Consider first this conditional claim: if all epistemic norms are equally unjustified, that the only *possible* form of justification is justification on the basis of operative norms. We need to be careful here; if relativism is false and absolutism about epistemic justification is true, then justification on the basis of operative norms isn't a *logically* possible form of justification, for the truth of absolutism logically excludes the truth of relativism. The only sense of possibility in which the conditional claim under discussion comes out true is *epistemic* possibility. Epistemic possibility is possibility conditioned upon what one knows. If what one knows is that *all epistemic norms are equally unjustified*, then two distinct options remain epistemic possibilities:

- (i) Epistemic justification is justification on the basis of (local) operative norms (a corollary to epistemic relativism, on Sankey's definition.)<sup>42</sup>
- (ii) We are not justified in believing that any epistemic norms (e.g. on Sankey's showing, any criteria or rules that may be employed to justify a belief) are true (an implication of scepticism).

While it's true that if 'the norms operative within a particular context provide justification for beliefs formed within that context' then the relativist *is* in a position to say that 'epistemic justification is relative to locally operative norms', it's not true that the interim conclusion that all epistemic norms are equally unjustified recommends that 'epistemic justification is relative to locally operative norms' *any more than it recommends* the sceptical conclusion (ii). And this is because (at least from what Sankey's told us) 'epistemic justification is relative to locally operative norms' is supposed to follow from (i). But the intermediate conclusion Sankey takes himself to be entitled to favours (i) no more than it favours (ii). Clearly (i) is simply an epistemic possibility given the intermediate conclusion. But so is (ii).

I think it should be clear that the relativist is in a position to claim that epistemic justification is relative to locally operative norms *only if* in violation of the very kind of 'non-arbitrariness' principle that is *relied on* in Stage One of the argument, in order to set aside foundationalism

as a viable approach to vindicating the justification of epistemic norms (in the service of reaching his intermediate conclusion that all norms are equally unjustified). The dilemma Sankey's would-be relativist faces is thus this: once the intermediate conclusion is reached, the positive acceptance of relativism over scepticism amounts to a dogmatic adoption of one epistemic possibility over another, with nothing (associated with the acceptance of the intermediate conclusion) to positively *recommend* one possibility *over* another on any rational basis. If Sankey embraces relativism over scepticism, then he can't at the same time rely on (1) in (in Stage One) disregarding foundationalism as a viable way to vindicate a given epistemic norm as justified. Put simply, if foundationalism is unacceptably arbitrary, then so is accepting – with the intermediate conclusion as collateral – relativism over scepticism. Alternatively: by accepting relativism over scepticism on the basis of the intermediate conclusion, Sankey's no longer entitled to his intermediate conclusion that all norms are equally unjustified, as he's forfeited the rationale for not engaging with the foundationalist<sup>43</sup>.

### 3.3 Taking stock: three conclusions

I want to turn now to three observations that can be drawn from the analysis of where the Pyrrhonian problematic, understood as an argument for epistemic *relativism*, goes wrong. The first is a point about *equivocation*, the second about *equipollence*, and the third about *neutrality*. All three of these points are connected.

Firstly, the point about equivocation. In the previous section, I argued that it's not nearly as seamless a move as Sankey had suggested from the claim that all epistemic norms are equally unjustified to the epistemic relativist's conclusion that epistemic justification is relative to locally operative norms. Now, the reason I gave was that (in short) the 'intermediate conclusion' provides no positive support for the relativist *beyond* the support it provides for a competing sceptical conclusion. I want to now go a bit deeper with this and suggest *why* one might be, albeit mistakenly, lured to thinking that relativism is not merely (as scepticism is) a live option, if one takes as collateral the premise that all norms are relegated to equal standing, but instead (and much more strongly) a kind of inevitable result of this premise.

In short, I think the problem is a tendency to equivocate on the notion of an *epistemic norm* and in particular, what it is for an epistemic norm to be operant in a particular context. This point takes a bit of unpacking. To begin, consider that *beliefs* about which epistemic norms are true – *viz.*,

which criteria or rules are the correct ones – can be mistaken. (This is something everyone agrees on, except for perhaps the most extreme subjective forms of relativism.)<sup>44</sup> Accordingly, unless epistemic (cultural) relativism is *already* true, it doesn't follow from the fact that, say, some epistemic norm E is *believed* to be true in a culture, that justification for beliefs *really is* conferred, in that culture, in the way that E purports. So an argument *for* epistemic relativism obviously can't assume that an epistemic norm really is *operating* in a context just because everyone believes it is – if operating means – that justification for beliefs, in that culture, *really is* conferred in the way that E says justification is conferred.

But here's another sense in which an epistemic norm might be thought to operate. We might say that (in the Azande culture) the epistemic norm, 'Poison Oracle', operates:

*Poison Oracle:* We are justified in believing the deliverances of the poison oracle.

*Poison Oracle* operates in Azande culture, in the attenuated sense that most people in that culture believe that justification really is conferred in the way specified by *Poison Oracle*. But given that unless we already are assuming a radical subjective relativism, there is a clear difference between:

- (O<sub>1</sub>) beliefs about how justification is conferred to beliefs; and
- (O<sub>2</sub>) how justification is conferred to beliefs,

we can't just move (in an argument *for* epistemic relativism) from (O<sub>1</sub>)-facts to (O<sub>2</sub>)-facts, where (O<sub>2</sub>)-facts depend on how epistemic norms *operate* in the robust sense of actually *holding* (and being such that: beliefs *that* they are true could aspire to correctness *regardless* of whether epistemic relativism is true). An equivocation on what it is for norms to be operant in a local context – that is, an equivocation between the attenuated and robust senses of what it is for an epistemic norm to be operant (O<sub>1</sub>) and (O<sub>2</sub>) – would suffice to explain why one might try to move from where Sankey took himself to be, as having established that all epistemic norms are on equal standing – *viz.*, a position which trades on an (O<sub>1</sub>)-fact about operant norms – *viz.*, that all *beliefs about epistemic norms are on the same epistemic standing*) to the relativist conclusion, which depends on an (O<sub>2</sub>)-fact about operant norms. In subsequent chapters, it will be important to keep in mind the ease by which this equivocation on the notion of an operant norm can occur, as it occurs often.

So much for the point about equivocation. Now for the point about *equipollence*. Agrippa's 'trilemma' – the key wheel of the Pyrrhonian regress-concerns three 'modes': infinite regress, hypothesis and circularity. But Agrippa had *five* modes. The other two are disagreement and relativity. Following Lammenranta (2008, 16), it's helpful to think of these five modes as 'working together' in the following respect:

disagreement and relativity challenge us to justify our beliefs, and then the rest of the modes show that the process of justification cannot be completed in a satisfactory way.

The sceptic does not, thus, 'run Agrippa's trilemma' (i.e. the 'justification-eater' function) on *every* single proposition considered. Only some of them. As Gail Fine (2000, 221) puts it:

Sextus claims that although skeptics lack beliefs about anything *unclear*, they do have some beliefs. In particular, they have beliefs about how they are appeared to: when it appears to them that honey sweetens, for example, they believe that it appears to them that honey sweetens. However, they suspend judgment as to whether the appearance is true; they do not believe (or disbelieve) that honey is sweet.

Fine is here parting ways with Descartes, who thought the ancient sceptics did not trust their non-doxastic appearances. Fine goes on to quote Sextus as saying:

Those who think the skeptics reject what is apparent have not, I think, listened to what we have to say. (PH 1.19, *ibid.* 221)

The matter of just *which* beliefs we can attribute to the Pyrrhonian is (as suggested at the beginning of this chapter) a matter of scholarly debate.<sup>45</sup> But it is clear enough that at least one safe dividing line – one which fits with Agrippa's five modes working together – is a dividing line between (as Fine attributes to Sextus) those things that are *unclear* and those things that are not. Now the mode of disagreement can make an issue unclear by presenting one with opposition, whereby one finds oneself in a position where a justification is requested, or would seem otherwise required, to rationally maintain one's previous position. One is not brought into equipollence willy nilly.

The reason this is important in the present context is because a question with both empirical and philosophical import is this: *if*, as the

Pyrrhonian suggests, the encountering of opposition – *viz.*, disagreement – is itself an *epistemically significant fact* – one which can engender epistemic demands for justification – then what kinds of actual and possible disagreements *about epistemic norms* (or, more generally, epistemic facts) are among the epistemically relevant ‘oppositions’ which one actually (and could possibly) encounter?

When Agrippa’s five modes are working in unison,<sup>46</sup> the kind of disagreement that engenders confusion and a corresponding obligation<sup>47</sup> to provide justification (one which leads one into Agrippa’s Trilemma) was a somewhat loaded concept. Such disagreements were characteristically regarded as ones where one finds equal arguments for one side as for another, a position (roughly) as convincing as one’s own. As Jonathan Barnes (1998, 59) puts it:

A Pyrrhonist asks ‘Is it the case that *P*?’ (‘Do there exist gods?’, ‘Can we discern true from false appearances?’, ‘Is the world a structure of atoms and void?’). He then assembles arguments in favour of an affirmative answer, and arguments in favour of a negative answer. The two sets of arguments exactly balance one another [withholding] supervenes, directed toward the proposition that *P*.

Contrast now the kind of cases Barnes envisions with a case where two very different epistemic systems collide (we’ll consider more such cases in Chapters 4 and 5). The Pyrrhonist suggestion that encounterings of such oppositions are epistemically significant for one who previously holds a viewpoint is a suggestion that is clearly embraced by a range of proponents of epistemic relativism – and it is very much to be found in the background of Richard Rorty’s (1979, 328–9) famous case that pits against each other Galileo and Cardinal Bellarmine.

Awkwardly, though, most cases in which the most radically divergent epistemic systems clash are ones where neither party is *in fact* (as a psychological point) drawn in to legitimate confusion. As Wittgenstein (OC §611) says, more typically, each calls the other a fool or a heretic. But even more, it’s not clear (in the extreme cases, cases often used to *motivate* epistemic relativism) how it is even supposed to work for one side to regard the other to be in a position such that, as Barnes characterises, one regards the case on each side to be equally balanced. Consider a quick example: The Azande provide 10 reasons to believe in the deliverances of the Poison Oracle; Richard Dawkins provides 10 reasons not to believe these deliverances. The number of reasons is balanced, neither side considers there to be equal reasons on each side. Neither feels the slightest bit confused or uneasy.<sup>48</sup>

Bringing this full circle: there seems to be both a (rather significant) psychological and epistemic difference between the disagreements that the Pyrrhonian regards as epistemically significant, and the kind of disagreements typically appealed to in the service of motivating epistemic relativism, where distinct systems clash. The difference is that in epistemic-system-clash cases (e.g. Dawkins and the Azande), and unlike in Barnes-style cases, there is neither actual confusion, nor does it seem clear how reasons (relative to each system) should be regarded by either side as even remotely balanced. A point of curiosity, which we'll take up in the next chapter (on non-neutrality and question-begging arguments for epistemic relativism), is *what the alleged epistemic significance of such disagreements would be explained by*.

Finally, I want to draw attention to a point about scepticism, relativism and *neutrality* that is conveniently a kind of 'bridge' between this chapter and the next. The version of the Pyrrhonian Problematic that Sankey appealed to was the *Regress* version. It is of course the most famous version of the problematic. But it's not the only one. As Lammenranta (2008) sees it, the most philosophically potent form of the problematic – what he calls the *dialectical version* – actually does not make any needed recourse to the regress that is generated by the three most famous of Agrippa's modes. It is instead centred around (and entirely so) the mode of *disagreement*; how, though, might one move from disagreement to scepticism *without* traversing Agrippa's modes? Here's Sextus:<sup>49</sup>

For we shall not be able ourselves to decide between our own appearances and those of other animals, being ourselves a part of the dispute and for that reason more in need of someone to decide than ourselves be able to judge. When the self-satisfied Dogmatists say that they themselves should be preferred to other humans in judging things, we know that their claim is absurd. For they are themselves a part of the dispute, and if it is by preferring themselves that they judge what is apparent, *then by entrusting the judging to themselves they are taking for granted the matter being investigated before beginning the judging*. (Sextus PH, *op. cit.*, 14)

On this passage, Lammenranta (2008, 13) remarks:

Here Sextus clearly thinks the mistake the dogmatists make is a dialectical one. When dogmatists judge that their own appearances are true while those of other animals and other people are false, they simply assume what they are supposed to prove. They beg the questions against their opponent's conflicting judgments. *So our inability*

*to decide between conflicting appearances is not a psychological matter but an inability to do so without violating the rules of dialectic. It is an inability to resolve disagreements without begging the question at issue.*

The italicised portion of the passage (my italics) is a reason Lammenranta takes to motivate (*via* the Pyrrhonian Problematic) a *sceptical* conclusion by way of motivating the premise that disagreements can be irresolvable due to the inescapable fact that one party is relegated to begging the question against the other. Interestingly, this very issue has in fact been appealed to by (among others) Boghossian (2006), Rorty (1979), Siegel (2011), Pritchard (2010), Hales (2014) (and many others) as a stock premise in an argument for epistemic *relativism*.

That is, the very *fact* that we can't (or so it seems) non-question-beggingly resolve disagreements about what epistemic principles, norms and facts are true has been famously regarded as a motivating reason for embracing the picture offered by the epistemic relativist, where justification is essentially local. (And so it might be true after all that there is a Pyrrhonian argument for epistemic relativism, then – just not the one we expected.) Regardless, we'll engage head-on with the *non-neutrality* argument for epistemic relativism in Chapter 4.

# 4

## Dialogic Arguments for Epistemic Relativism

*Abstract.* Galileo and Cardinal Bellarmine couldn't agree about the truth of Copernican heliocentricism. But, as Richard Rorty (1979) famously highlighted, they also couldn't agree about what evidential standards were even relevant to settling the matter. The inability of interlocutors to non-question-beggingly break the deadlock in cases like this – where there appears to be a deep clash at the level of epistemic systems – has led some philosophers to think that the *only* sense in which either party can be correct is relative to their own epistemic system. The present aim will be to critically engage with arguments for epistemic relativism that take this general shape – what I call *dialogic arguments* – which point to certain properties of actual (or possible) dialogues and conclude on the basis of the presence of these properties that epistemic relativism is true. Dialogic arguments can be 'actualist' or 'possibilist', depending on whether the dialogues meant to be doing the relevant work are regarded as actual. I argue on empirical grounds, with reference to the literature on cognitive biases, that actual dialogues are ill-suited to motivating epistemic relativism. I conclude by suggesting why retreating to a 'possibilist' strategy is not promising; finally, I show that even if the problems I've raised can be overcome, dialogic arguments leave us no closer to epistemic relativism than to scepticism.

### 4.1 A final card for the Pyrrhonian?

The previous chapter did not leave things on a very optimistic note, at least, in so far as the Pyrrhonian Problematic, featuring the famous Agrippan Trilemma, was meant to underwrite a compelling argument

for epistemic relativism. But recall that, as we left things, there was a Pyrrhonian-style card left to play.

Agrippa's modes of *hypothesis*, *circularity* and *infinite regress* are the modes most commonly associated with the Pyrrhonian challenge. But they are meant to work in tandem with the mode of *disagreement*. Maybe disagreement can do some important work in its own right. Here's a simple picture: the mode of disagreement is 'up first':<sup>1</sup> disagreement reveals that there are competing claims on the matter of *p*, appreciation of which challenges us to justify our belief that *p*; once we accept the challenge, the rest of the modes show that the process of justification cannot be completed in a satisfactory way.<sup>2</sup>

Lammenranta (2008, 16) refers to this line of thinking as the *dialectical interpretation* of the Pyrrhonian problematic, which he regards as the most promising picture of how the modes work together in the service of motivating scepticism. To more clearly appreciate this, just suppose there is a dispute about whether *p*:

*Template Dialectical Argument for Scepticism*

1. *A* believes that *p*.
2. *B* believes that  $\sim p$ .
3. At most one of them is right.
4. The disagreement between *A* and *B* is irresolvable.
5. We should suspend judgment about *p* (scepticism)

Importantly, on the dialectical interpretation, the more 'famous' Agrippan modes (hypothesis, circularity and infinite regress) come in to play only once one attempts to reject premise (4). And so, if the Agrippan modes do their work, they do it in preventing one from avoiding the conclusion (5) of the above argument by denying premise (4). We can see then that, as Lammenranta (2008, 16) puts it, when we set things up this way, the mode of disagreement is 'the central one, and the other modes are subordinate to it'.

Though, in the face of the Template Dialectical Argument for Scepticism, there is an entirely different, and anti-sceptical, direction one might be inclined to take things. Some philosophers think that once we think we've got good reason to accept (1), (2) and (4), we should reject the sceptical conclusion by rejecting (3), the claim that *at most one of them is right*. Perhaps some of the considerations that lead to the disagreement's being irresolvable are evidence that we should reject (3) for a premise according to which *both parties are right*. And this 'relativist hijacking' of the Pyrrhonian's dialectical argument, involves – once (3)

is rejected – simply swapping out the sceptic’s conclusion (5) for the relativist’s conclusion, (6), as follows:

Template Dialectical Argument for Scepticism (Relativist’s Twist)

1. *A* believes that *p*.
2. *B* believes that  $\sim p$ .
3. ~~At most one of them is right.~~
4. The disagreement between *A* and *B* is irresolvable.
5. ~~We should suspend judgment about *p*.~~
6. *A* and *B* are both right; *p* is true relative to *A*’s perspective,  $\sim p$  is true, relative to *B*’s perspective. (*Epistemic relativism*)

And indeed, one of the most popular contemporary strategies for motivating epistemic relativism proceeds in just this way, where the ‘irresolvability’ at play in (4) is typically claimed as a feature of ‘epistemic systems clash’ cases, where there is (i) a first-order disagreement about some target proposition, and then (ii) a second-order disagreement about what kinds of epistemic norms are germane to settling the disagreement about the target proposition.

The most famous such case in the relativist literature – thanks in large part to Richard Rorty (1979) and Paul Boghossian (2006a) – involves a 17th-century dispute between Catholic Cardinal Bellarmine and Galileo about Copernican heliocentrism. Bellarmine and Galileo disagreed about Copernican heliocentrism, but even more, they disagreed about what kinds of evidence are even relevant to settling this dispute.

Galileo had argued for the Copernican picture on the basis of telescopic evidence. Cardinal Bellarmine was not convinced. By appeal to Scripture, he dismissed Galileo’s suggestion that Earth revolves around the sun as heretical. He noted further that:

It would be just as heretical to deny that Abraham had two sons and Jacob twelve, as it would be to deny the virgin birth of Christ, for both are declared by the Holy Ghost through the mouths of the prophets and apostles.<sup>3</sup>

From these antipodal perspectives, it looks like a quick route to impasse. Before digging in to this particular case and other such arguments, it’s helpful to clarify a point about the basic *structure* of such arguments. What *kind* of move, after all, is the relativist making when proposing a passage from (4) to (6) – from the datum that some disagreement has the property of being irresolvable to the conclusion that epistemic relativism is true?

## 4.2 The structure of dialogic arguments for epistemic relativism

Premise (4) in the template argument, the claim that the disagreement between *A* and *B* is irresolvable, is what we can call a *dialogue fact*, a fact that is purported to hold because of how some actual (or possible) dialogue goes. Call arguments that attempt to move from dialogue facts to the epistemic relativist's conclusion *dialogic arguments for epistemic relativism*:

**Dialogic arguments for epistemic relativism:** attempt to establish the epistemic relativist's conclusion by pointing to certain properties of actual (or perhaps also possible) dialogues and concluding on the basis of the presence of these properties that epistemic relativism is true.

These kinds of arguments will be the focus of this chapter. One rationale for taking the dialogic road to relativism that, in the face of dialogues that seem utterly deadlocked, relativism, as Steven Hales (2014, 63) has argued, provides a more compelling *resolution* – one on which ‘everyone wins’ – than do the most salient competing options. Those which he identifies are: keep arguing until capitulation, compromise, locate an ambiguity or accept scepticism<sup>4</sup>.

But what does it mean to say that relativism is a way to ‘resolve’ otherwise irreconcilable disagreements (2014, 63)? Let's take an example from Hales, originally owing to Putnam. How many things are pictured in Figure 4.1?



Figure 4.1 3 or 7?

You say 3. I say 7. We shout about this for a while. You think there are 3 because you embrace an atomistic mereology. I say there are 7 because I embrace unrestricted mereological composition (UMC). And then comes relativism to the rescue: I'm right, according to UMC, you're right according to atomism.<sup>5</sup> At this point, though, an important clarification is needed: a *relativist* resolution can't unpack 'I'm right according to X and you're right according to Y' as *merely* involving the embracing of conditionals such as: *If UMC, then 7 is right* and *If atomism, 3 is right*. After all, absolutists can accept these conditionals. Rather, a relativist resolution must involve some commitment to denying that there is *any further sense in which each can be right*.<sup>6</sup> For the relativist, there is only perspective/standard-relative correctness.

### 4.3 'No neutrality, therefore relativism'

Probably the most prevalent form of dialogic arguments are *non-neutrality arguments*: the slogan for such arguments is, following Harvey Siegel (2011, 206<sup>7</sup>): 'no neutrality, therefore relativism'.<sup>8</sup>

The particular dialogue fact that non-neutrality arguments point to is the following:

**Dialogic fact (non-neutrality):** in some actual disagreements where interlocutors disagree about epistemic standards (e.g. epistemic norms), there are no neutral 'meta-' or higher-order epistemic standards available to which we can appeal that will fairly or non-question-beggingly resolve our dispute.

Let's return now to the specific dispute between Galileo and Bellarmine, and see how this is supposed to work within the purview of the 'no-neutrality, therefore relativism' line of thinking. While Galileo and Bellarmine in fact disagreed about a range of things, the more narrowly defined dispute Siegel focuses on, in characterising the 'no neutrality, therefore relativism' line, is the existence of moons surrounding Jupiter – and what evidential standards are relevant to determining this. Siegel (2011, 205–6) understands the epistemic relativist's reasoning, with respect to what transpired between Galileo and Bellarmine, as taking the following shape:

Not only did the two parties disagree as to the truth of the relevant claim – Galileo affirmed the existence of the moons, while his opponents denied it – they also disagreed about the relevant standards (telescopic observation? naked eye observation? Scripture? Aristotle?)

to which appeal should be made in order to resolve their disagreement. *The relativist here claims that there can be no non-relative resolution of the dispute concerning the existence of the moons, precisely because there is no neutral, non-question-begging way to resolve the dispute concerning the standards.* Any proposed meta-standard that favors regarding naked eye observation, Scripture, or the writings of Aristotle as the relevant standard by which to evaluate “the moons exist” will be judged by Galileo as unfairly favoring his opponents since he thinks he has good reasons to reject the epistemic authority of all these proposed standards; likewise, any proposed meta-standard that favors Galileo’s preferred standard, telescopic observation, will be judged to be unfair by his opponents, who claim to have good reasons to reject that proposed standard. In this way, the absence of neutral (meta-) standards seems to make the case for relativism.<sup>9</sup>

Siegel’s relativist’s diagnosis of the dispute is one where the ‘non-neutrality’ that’s supposed to be motivating epistemic relativism is itself a matter of there being no appropriately neutral meta-standards which either could appeal to in order to resolve their dispute. And this is claimed to be a *result* of their disagreeing at the first- as well as the second-order.

It’s helpful here to distinguish between an appropriately *neutral* meta-standard and an appropriately neutral meta-standard *which either party could appeal to in order to resolve their dispute*. Presumably, a meta-standard will be appropriately neutral between the two parties to a dispute provided either could appeal to that meta-standard without begging the question against the other. As such, a meta-standard appropriately *neutral* to the debate in question is: ‘logic’. Or, even more specifically: *modus ponens*.

$$\frac{P \rightarrow Q, P}{\therefore Q}$$

Plausibly, Galileo and Bellarmine could both appeal to *modus ponens* in an attempt to adjudicate between whether the Bible or the telescope is better evidence about moons, without begging the question against one another.

But equally plausibly – and this is a point easily overlooked – this norm is simply *too* neutral. Imagine the following dialogue:

**Galileo:** P, because the telescope says so.

**Bellarmino:** Not-P, because the Bible says so.

**Galileo:** The Bible can't be as good a source of evidence about moons as the telescope (pointing to the telescope).

**Bellarmino:** The telescope can't be as good a source of evidence about moons as the Bible (pointing to the Bible).

**Galileo:** Well, I believe in *modus ponens* – so there's that!

**Bellarmino:** Me, too!

**Galileo:** Ok...great. That doesn't really help us, does it?

While neither would be in a position to object to each other relying on *modus ponens*, *modus ponens* is plausibly not going to be appropriately *discriminatory*<sup>10</sup> and this precisely *because* of its neutrality. (Likewise, suppose we substituted the equally unobjectionably neutral norm: *if A, then infer A* for *modus ponens*.) Unsurprisingly, there is another side to this coin: a meta-norm that does very well in the 'discriminatory role' will plausibly fail to be appropriately neutral.<sup>11</sup> (At one limit of discriminatoriness, the meta-norm *just* is the second-order norm.)

Putting these points together, a meta-standard can play the kind of role that it would need to play in order to bring interlocutors locked into an otherwise irreconcilable position into a non-questionbegging resolution, only if it is *both*: (i) *appropriately neutral*, such that it can be appealed to non-question-beggingly by either side; and (ii) *appropriately discriminatory*: not epistemically *inert*. Call an epistemic meta-norm that is appropriately neutral yet appropriately discriminatory *Archimedean*.<sup>12</sup>

We can now spell out the 'No-neutrality, therefore relativism' argument (with reference to the Galileo/Bellarmino dispute as follows):

*No-neutrality, Therefore Relativism*

7. There can be a non-relative resolution of the dispute concerning the existence of the moons, only if there is an *Archimedean meta-norm available* (i.e. *appropriately neutral* and *appropriately discriminatory*)
8. In the context of the dispute between Galileo and Bellarmino, no such Archimedean meta-norm is available.
9. Therefore, it's not the case that there can be a non-relative resolution of the dispute concerning the existence of the moons.
10. Therefore, epistemic relativism is true.

Obviously, (9) doesn't entail (10). The relativist needs some further 'bridge' premise. But let's grant (for now) that establishing (9) – that it's *not* the case that there can be a non-relative resolution of the disagreement in question – at least brings one closer to epistemic

relativism. I want to now turn to why we should be sceptical about (8) and, thus, (9).

Consider the rationale that is being offered for (8), the claim that in the context of the dispute between Galileo and Bellarmine, no Archimedean meta-norm is available. The *rationale* for (8), offered in the quoted passage from Siegel, identifying the (no-neutrality-therefore-relativism) relativist's reasoning, looks like this: that it is *because* Galileo and Bellarmine disagreed at the first order, *and* at the second order about what epistemic standards are best suited to adjudicating the first-order dispute, that therefore no Archimedean meta-norm is available. Thus, in short: first order + second order disagreement → no Archimedean meta-norm is available.

I want to turn now to an argument from parity in order to highlight why this rationale (e.g. the first- order and second- order disagreement rationale) for securing (8) doesn't hold water.

#### 4.4 Argument from parity

Let's move now from a 17th-century debate about the position of celestial bodies to a contemporary debate about what is involved in knowing how to do something. The primary dividing line in the contemporary know-how debate pits on opposing sides *intellectualists* and *anti-intellectualists*. The quick and easy picture is this: intellectualists claim that knowing how to do something is in virtue of *propositional knowledge*. By contrast, *anti-intellectualists* claim that knowing how to do something is in virtue of possessing *abilities*, rather than knowing propositions.<sup>13</sup>

Interestingly, these positions have, at least *de facto*, lined up with second-order positions about the kind of evidence that is most relevant to establishing what it is in virtue of which one knows how to do something. Typically, philosophers who take a stand on the intellectualism/anti-intellectualism debate primarily by appealing to linguistic evidence have supported intellectualism. Likewise, philosophers who take a stand on intellectualism versus anti-intellectualism primarily by appealing to evidence from cognitive science have supported anti-intellectualism.<sup>14</sup>

Against this very brief background, I want to consider, as a case study, a very *specific* debate between intellectualist Jason Stanley (2011) and anti-intellectualist Josefa Toribio (2008). Stanley and Toribio disagree about how to think about this (somewhat bizarre) case of patient 'DF' (Goodale and Milner (1992)), a case famous in visuomotor research.

**Case of DF:** DF had a form of brain damage that, due to carbon-monoxide poisoning that caused bilateral damage along the ventral stream of visual processing, meant that she was left unable to recognise the ‘size, shape and orientation of visual objects’. DF could, nonetheless (given the intact character of her *dorsal* stream of visual perception) retain accurate ‘guidance of hand and finger movements directed at the very same objects’ (*ibid.*). DF can, for instance, reliably place a letter through a rectangular slot, though prior to reaching out, DF does not know what the orientation of the slot is.

Interestingly, Stanley and Toribio have very different ideas about what the DF case indicates about the truth of intellectualism. Consider the following ‘target’ proposition:

**Target proposition (P):** The DF case counts *against* intellectualism.

Toribio thinks P is true; Stanley thinks P is false. Think of this as their ‘first-order’ disagreement. Furthermore, Stanley’s and Toribio’s respective rationales for their conflicting first-order positions seem heavily informed by what kind of evidence they regard as most appropriate to *evaluating* P. (Think of this as their *second-order* disagreement.)

Toribio thinks that the case of DF pretty clearly shows that the intellectualist line must be wrong, since it looks like DF knows how to put the letter through the slot, even though her doing so couldn’t possibly (and *contra* intellectualism) be guided by her propositional knowledge of the way to do this, because, claims Toribio (2008, 43–44):

[DF] cannot perceive the features, e.g. the orientation, that govern her motor behavior in the posting task, and hence couldn’t recognize them as in any way constituting a reason for her action. DF lacks the kind of phenomenal experience that would underwrite an appreciation of her own behaviour as suited to solving the problem.<sup>15</sup>

Stanley (2011, 172), who already takes the *linguistic* evidence for intellectualism to be compelling (e.g. 2011), does not regard the cognitive–scientific evidence Toribio points to as of a sort that would count against intellectualism. Here are his remarks:

Suppose we take Goodale and Milner’s results to show that DF does not know what the orientation of the slot is. That does not entail that DF’s action is not guided by knowledge how, in my sense. At most it

would show that the possession of the knowledge how to fit cards into slots, in my favored sense, does not require knowledge of what the orientation of the slot is. In short, *at most* what DF shows is *that one can know how to post a card into a slot without knowing what the orientation of the slot is*.<sup>16</sup>

Stanley adds, shortly after, that:

*All it shows* is that one can have propositional knowledge concerning a way of putting a card into a slot, without knowing the orientation of that slot. DF shows that propositional knowledge concerning a way of putting a card into a slot, *contra* Toribio, does not require perceiving all the features of that way. (*Ibid.*, 172, my italics)

Notice that Stanley and Toribio are disagreeing about not only the first-order matter of *whether* the case of DF counts against intellectualism. Their disagreement seems to run deeper. Put simply: Stanley, by interpreting the cognitive scientific evidence so as to be compatible with a view he already thinks is compelling on linguistic grounds, is at least tacitly giving a kind of priority of relevance to the linguistic evidence.<sup>17</sup> Toribio by contrast is giving a priority of relevance to the cognitive scientific evidence. She thinks cognitive scientific evidence is good grounds for rejecting intellectualism *even if* Stanley is right that the linguistic evidence counts in favour of it. What we have here is, as with the case of Rorty and Bellarmine, both first- *and* second- order disagreement.

Recall now: the *rationale* for premise (8) in the 'no neutrality, therefore relativism' argument adduced by Siegel's relativist was supposed to be this: *first order + second order disagreement* → *no Archimedean meta-norm is available*. The thrust of the argument from parity is that Stanley v. Toribio is a case that very plausibly features both first- *and* second- order disagreement, *but*, and contrary to the rationale behind (8), it's *not* the case that *there could be no Archimedean meta-norm available*.

There is a neat and tidy way to make this point, especially given that the know-how debate is a special case where there are really *three* central varieties of evidence on the basis of which philosophers have taken a *stand* on whether intellectualism is true. There is linguistic, cognitive-scientific, and also phenomenological evidence<sup>18</sup>. Hubert Dreyfus (2005), for example, is an important player in the know-how debate,<sup>19</sup> and it's a debate he's entered entirely on the back of his work on the phenomenology of skilled action. His primary source of appeal is Heidegger.<sup>20</sup>

*Question:* could phenomenological evidence, with respect to the matter of 'P' (the matter of whether DF counts against intellectualism) potentially be both (i) *appropriately non-questionbegging* in the context of the dispute between Stanley and Toribio; and (ii) *appropriately discriminatory*?

The answer to (i) is yes. This is not to suggest that either Stanley and Toribio would in fact be inclined to appeal to phenomenology. Question (i) gets a 'yes' answer so long as it's true that *if* either did, it needn't be question-begging. Phenomenological evidence, as such, is evidence one could in principle appeal to *regardless* of whether one already holds some antecedent view about the comparative import of cognitive–scientific as opposed to linguistic evidence with respect to P. Moreover, (ii) gets an affirmative answer. Phenomenological evidence *could* be discriminatory (unlike, say, a maximally neutral but non-discriminatory meta-norm such as: infer *a* from *a*).

Putting this all together, Toribio and Stanley disagree at the first- and second- order, *vis-à-vis* P. It's *not* the case that, in the context of their debate, there simply *can be no Archimedean meta-norm*; for everything that's been said, phenomenological evidence *could* play such a role – *viz.*, phenomenological evidence could be both appropriately neutral *and* appropriately discriminatory. Therefore, it is false that, *in virtue of* first- and second- order disagreement, there can be no Archimedean meta-norm available.

But this conclusion undercuts the support *originally* proposed for (8), as per Siegel's envisioned relativist. After all, premise (8) of the 'No-Neutrality, Therefore Relativism' argument claims that, in the context of the dispute between Bellarmine and Galileo, no Archimedean meta-norm is available. But *this* (at least, as Siegel had painted the picture on behalf of the relativist) was supposed to be precisely *because* Bellarmine and Galileo disagreed at both the first- and second- orders. Stanley v. Toribio shows that if Galileo and Bellarmine are in a position such that there is no Archimedean meta-norm available, it will not be simply in virtue of their disagreeing at the first- and second- order.

#### 4.5 Strengthening the argument: Hales

The upshot from the parity argument is that first- and second- order disagreement simply isn't enough to secure the unavailability of an Archimedean meta-norm.. But perhaps some dialogues (perhaps, even the Bellarmine–Galileo dialogue, for all we've said) have some other *further properties* in virtue of which there *really could be no Archimedean meta-norm*.

On this note, consider Steven Hales' (2014, 78–80) latest case, featuring 'Jack' and 'Diane':

Jack and Diane [...] disagree over P: human beings each have a soul which animates their bodies and is immortal [...] Jack denies P and Diane affirms it. Jack and Diane further disagree about what kind of evidence is relevant to settling their dispute. Jack maintains that the appropriate evidence is provided by the analytic rationalist methodology of contemporary philosophy of mind, including reflection on hypothetical cases, thought experiments, and appeals to intuition. Zombies, swampmen, Chinese rooms, C-fibers, and strangely reared neuroscientists figure prominently in Jack's reasoning. Diane avers that the appropriate evidence is provided by the Bible, along with its interpretation by the learned doctors of the church. Jack reports Jaegwon Kim's observation that there is an almost complete consensus among philosophers in rejecting the existence of an immortal spiritual soul, Diane quotes *The Catechism of the Catholic Church*. (Pt 1, sec 2, ch 1, art 1, para 6, §366)

As an initial observation, Hales' Jack and Diane look to be in the same position as Galileo and Bellarmine, and for that matter, Stanley and Toribio: we have first- and second- order disagreement – disagreement about a target proposition and disagreement about the evidential standards germane to the target proposition. And it's already been shown that this doesn't simply entail the unavailability of an Archimedean meta-norm. But, interestingly, Hales' case continues:

In the present example, Jack and Diane have a genuine irreconcilable difference; they disagree over proposition 'P', they disagree over what evidence is relevant to establishing to truth or falsity of 'P', and they have no additional means of settling their debate about the relevant evidence. Jack and Diane *cannot* discover any mutually agreeable meta-evidence which would allow them to settle their dispute over first-order evidence...we might regard a persistent failure to agree about even the meta-evidence for a claim as a good reason to conclude that *there is no such thing as the right kind of first-order evidence*. In such a case, provided we are not tempted by scepticism, relativism appears to be our last option. The dispute between Jack and Diane is resolved by determining that 'P' is both true and false. 'P' is true relative to Diane's perspective, a perspective which includes as an epistemological component the methodology of appeal to revelation, the

Bible, and its expert interpreters as a source of noninferential beliefs. 'P' is false relative to Jack's perspective, the epistemology of which includes analytic rationalism.<sup>21</sup>

Hales' case is a clever one. Notice that this case bypasses the issue of what would be needed to secure the unavailability of an Archimedean meta-norm. In this case, that there is an inability to locate an Archimedean meta-norm is simply *built in* to the details of case.

Consider now a revamped version the original non-neutrality argument, one which makes recourse to the particulars of Hales' envisioned Jack & Diane dialogue:

*Revamped Non-Neutrality Argument for Epistemic Relativism* (Hales' Variation)

11. There can be a non-relative resolution of the dispute between Jack and Diane concerning the soul, only if there is an *Archimedean meta-norm* (e.g. *appropriately neutral* and *appropriately discriminatory*) available.
12. In the context of the dispute between Jack and Diane, no such Archimedean meta-norm is available (*ex hypothesi*).
13. Therefore, it's not the case that there can be a non-relative resolution of the dispute between Jack and Diane concerning the soul.
14. Therefore, epistemic relativism is true.

Two points to note about the Hales-style version of the argument, which moves from Jack and Diane dialogue facts to epistemic relativism. The first is that Hales (just as anyone else defending this kind of argument) needs a 'bridge premise' to get him from (13) to (14), because (14) certainly doesn't follow from (13).

And to his credit, Hales does positively defend his move from (13) to (14), *via* a kind of 'process of elimination strategy'. As he sees it, in a situation like Jack and Diane find themselves in (one where Hales' description entails there can be no Archimedean meta-norm available), the most salient available options are: (i) keep arguing until capitulation, (ii) compromise, (iii) locate an ambiguity or contextual factors, (iv) accept scepticism, or (v) adopt relativism.

Hales himself thinks (i)–(iii) are not very promising, and while he says (iv) can't be ruled out as a viable option, he regards it as 'throwing in the towel' in a way that relativism is not. (We'll return to this.) As with the previous version of the argument considered, let's assume for now that if Hales can get to (13) – the claim that it's not the case that there

can be a non-relative resolution of the dispute between Jack and Diane concerning the soul – then he is closer to (14) than if (13) were not established.

The second observation of the Hales variation on the argument is this: Hales has achieved (13) – the premise that that there *can* be no available Archimedean meta-norm, for Jack and Diane, by *stipulation*. Again, part of the details of his case, after all, is that Jack and Diane’s position is *irreconcilable*. More specifically, he says that:

Jack and Diane *cannot* discover any mutually agreeable meta-evidence which would allow them to settle their dispute over first-order evidence.

If there *were* available an Archimedean meta-norm, then they *de facto* wouldn’t be in the irreconcilable position that Hales stipulates they’re in. *But wait*: is he allowed to just stipulate something that entails there could be no Archimedean meta-norm available? Of course he is. But, as we’ll soon see, there might be an eventual price for this. As a matter of fact, dialogic arguments for epistemic relativism have typically regarded *actual* dialogue facts as supporting a stronger case for epistemic relativism than merely *possible* dialogue facts.<sup>22</sup> At any rate, Hales has two options: he can *either* argue that epistemic relativism is motivated by *possible dialogue facts* or rely on *actual dialogue facts*. The obvious question now is: *does anything prevent Hales from supposing there are actual dialogues that could do the work he’s using the Jack and Diane case to do, to support epistemic relativism?*

In what follows, I’ll argue on empirical grounds, that there is good reason to deny that *actual* cases can play the kind of role that would be needed to motivate epistemic relativism. And if the empirical argument is successful, then a Hales’-style proponent is forced to retreat to a defence of *possible* dialogue facts as what’s motivating epistemic relativism. I then conclude by raising some problems for a possibilist strategy and showing that even if these problems can be overcome (and premise (13) can be defended), we are left ultimately no closer to relativism than to scepticism.

#### 4.6 Cognitive Bias: Some Highlights

What judgments I make (in a dialogue or otherwise) are to some extent influenced by what epistemic principles I accept – e.g., whether I accept epistemic principles like the ones Galileo did or like the ones Bellarmine

did. But what epistemic principles I accept aren't the only things that influence what kinds of judgments I make, and the kinds of shapes that my interactions with my interlocutors take. Recent empirical psychology has shown that, across a range of well-documented cases, what judgments we make are in fact rife with cognitive biases of various kinds, most of which we are completely unaware. I want to turn to a quick summary of some of the highlights of this recent literature, after which I'll suggest why the extent to which biases affect our judgment is in a crucial way relevant to the viability of arguments to the effect that irreconcilable arguments are evidence that the truth of the contested judgments should be relativised to the differing epistemic principles.

#### 4.6.1 *Situationism*

For one thing, our judgments are highly sensitive to *situational factors*, features of our immediate environment, entirely irrelevant to the issue at hand, and – importantly – *unconnected to which epistemic principles we embrace*. John Turri (2015) offers a helpful overview of some well-studied examples of such situational factors that influence judgment.

We're less likely to recognize someone's face after working on difficult crossword puzzles than reading; we overestimate distances and upward angles when tired or carrying heavy equipment; we're worse at judging distances in hallways than in a field; we're more likely to accept a written claim as true when it's easy to read; we're more likely to judge someone credible who speaks quickly; we're more likely to think that easy to pronounce stocks will outperform difficult to pronounce ones.

Turri 2015, §2

This is really just the tip of the iceberg. But what does this *mean*? Some quick background will be helpful. The relevance of these kinds of situational factors to human performance, more generally, was originally drawn out in detail in the ethics, rather than the epistemology, literature – by Doris (1999; 2000) and Harman (2002) – in order to challenge the empirical adequacy of *virtue ethics*. The challenge, in short, is that, given the surprising extent to which situational factors (e.g. the presence of bystanders, mood elevators, ambient sounds and such) are actually shown to influence morally relevant behaviour, the stable character traits postulated by virtue ethicists can't be doing the explanatory work that, in theory, they are supposed to be doing. That is, our dispositions to

morally relevant behaviour lack the kind of cross-situational consistency that matches the description of virtues. As Alfano and Loeb (2014 §3.1) put it, describing Doris's (2002) strand of this challenge:

the best explanation of this lack of cross-situational consistency is that the great majority of people have local, rather than global, traits: they are not honest, courageous, or greedy, but they may be honest-while-in-a-good-mood, courageous-while-sailing-in-rough-weather-with-friends, and greedy-unless-watched-by-fellow-parishioners.

The empirical argument from moral psychology to the conclusion that virtue ethics is inadequate has come to be called the *situationist challenge* to virtue ethics.<sup>23</sup> It's not a stretch to envision how this argument strategy can be extended from ethics to epistemology; after all, the kinds of examples noted in Turri's 'highlights' quote above indicate that, regardless of what situational factors affect morally relevant behaviour, they also affect our *cognitive* performance. Thus, in light of the situationist challenge to virtue ethics, it should really be no surprise that attention to the extent to which situational factors influence performance in cognitive tasks stands to challenge *virtue epistemology*, which submits a claim that roughly parallels the claim of the virtue ethicist: that intellectual virtues or abilities play a significant role in explaining *cognitive* success.

In recent work, Mark Alfano<sup>24</sup> has led the charge on this score, tracing out implications of a range of biases in cognitive task performance for epistemology. And Alfano's charge, suitably understood, needn't be constrained to virtue epistemology.<sup>25</sup> In fact, the charge pared to its simplest form is a problem for any view in epistemology (virtue-theoretic or otherwise) which relies in one way or another on the reliability of our inductive reasoning abilities.

Drawing from work by Kahneman and Tversky, Alfano draws attention in particular to the *availability* and *representativeness* heuristics. The availability heuristic (e.g. Tversky and Kahneman (1973)) leads people to expect that the probability of an event or proportion of some property, in a population is positively correlated with the ease by which the event or property in question can be brought to memory. A simple study that illustrates this idea is one conducted by Tversky and Kahneman (1973), in which participants were asked: 'If a random word is taken from an English text, is it more likely that the word starts with a 'K' or that 'K' is the third letter?' In fact, it approximately twice

as likely a given word has 'K' as the third rather than first letter; however, given that it is typically easier to think of words with 'K' in the first, rather than third letter place, participants significantly assessed the probability as higher that 'K' is the first letter than the third.<sup>26</sup> One factor that is positively correlated with ease of recall (which then influences probability estimates) is how *recently* one has been exposed to some idea or concept. In a study by deTurck et al. (1990), for example, mock jurors were significantly more inclined to regard a testifier as deceptive if the testifier lied after telling the truth, than if the testifier lied and then told the truth; more generally, information encoded from our most recent encounters typically is (unconsciously) afforded disproportionate weight.

Another well-studied cognitive bias Alfano draws attention to is the *representativeness bias* (Tversky and Kahneman (1974)). Consider here the case of Linda. 'Linda is thirty-one years old, single, outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.' Participants, in light of this information, were asked to assess the probability that various attributes are true of Linda. Among them were that (i) Linda is a bank teller; and (ii) Linda is a bank teller and is active in the feminist movement. It was found that 89% of participants rated (ii) more probable than (i), even though this is probabilistically impossible.<sup>27</sup> What explains this mistake according to Tversky and Kahneman (1974) is our instinctive use of heuristics, such as 'stereotyping', or judging likelihood/frequency of something on the basis of its perceived resemblance to the stereotype of the item or individual in question.

#### 4.6.2 Implicit bias

A particular strand of cognitive bias that turns out to be especially prevalent when we interact with other people – as we do when engaging in dialogue – is *implicit bias*. Implicit bias consists in our 'unconscious tendencies to automatically associate concepts with one another' (Saul 2013, 244), and in particular, when we associate implicitly certain context-specific performance behaviours (e.g. academic performance, athletic performance, intelligence) with concepts such as racial category, religion, etc.

As Saul (2013, 244) notes, implicit biases can and often do lead to some 'disturbing errors'. One nice and clean example she offers involves our perception of CV quality. Saul notes that, in cases where

the experimenter holds fixed all items on a CV, switching out only the names at the top, what is found is that:

[T]he same CV is considered much better when it has a typically white rather than typically black name, a typically Swedish rather than typically Arab name, a typically male rather than typically female name, and so on.<sup>28</sup>

Such judgments are clearly being influenced by something which, as Saul notes, should be entirely *irrelevant*, which is social category. Another striking example of implicit bias Saul draws attention to is ‘shooter bias’, which skews perception, as in, what objects we think we are seeing. ‘Shooter-bias’ cases (e.g. Correll et al. (2006); Greenwald et al. (2003)) are well-studied and reveal that a given ambiguous object is significantly more likely to be perceived as a gun when held by a young black or, as Unkelbach et al. (2008) have shown, a young Muslim man, and as something ‘innocent’ (e.g. like a phone) when held like by a young white man.<sup>29</sup>

Obviously, the extent to which studies like the above should concern epistemologists who want to maintain a purer picture of our epistemic practices (where cognitive success is explained by our cognitive abilities and our adhering to various epistemic principles) depends on just how prevalent these kinds of biases really are. If they were reasonably rare, for instance, the damage to the received picture might be limited.

However, such biases are not at all rare, and they are especially prevalent when our judgments are in some relevant way influenced by our *interactions with other individuals*. Saul (2012, 245–146) writes:

The problem starts to become vivid when we ask ourselves when we should be worried about implicit bias influencing our judgments. The answer is that we should be worried about it whenever we consider a claim, an argument, a suggestion, a question, etc. from a person whose apparent social group we’re in a position to recognize. Whenever that’s the case, there will be room for our unconscious biases to perniciously affect us.

#### 4.7 Cognitive Bias and Epistemic Relativism

The take-home lesson from the previous section’s highlights can be summarised as follows. Our cognitive lives are riddled with biases, and this is well-documented with respect to how we: (i) are influenced by

situational factors in the environment around us; (ii) are subject to a range of heuristics and biases (e.g. availability, representativeness, etc.) which influence to a significant extent the kinds of inferences we make; and (iii) are subject to implicit biases that infiltrate to a significant extent our interactions with our interlocutors. While implicit biases directly threaten to influence the shape our debates take with our interlocutors, (i) and (ii) are often at work as well, as our exchanges with other individuals are not insulated from either (i) situational factors present in our local environment, or (ii) any of the many well-documented cognitive heuristics and biases that we so often employ. And here's the key fact: *none of these factors influencing the shape our debates take are factors that are the product of which epistemic principles we endorse.*

I now want to build upon this point, and defend the following: the pervasiveness of our biases, as they manifest in our exchanges with our interlocutors, calls into doubt that *actual* dialogue cases are germane to playing the kind of role that they would need to play to plausibly motivate epistemic relativism. And, if *this* is right, then the upshot is that Hales-style cases will have to be argued to motivate epistemic relativism *qua* mere *possible* disagreements – the ramifications of which will be evaluated.<sup>30</sup>

But first, the task is to defend the claim that the prevalence of our cognitive biases calls into doubt that actual cases can play the kind of role that they would be needed to play to motivate epistemic relativism. In order to make this case, I think it will be helpful to reflect on a thought experiment where the irreconcilability of the dialectical position is 'triggered' by features of the environment or other biases.

#### 4.7.1 Triggered disagreements

Let's compare now Jack and Diane's irreconcilable disagreement, with another:

**John and Lise** (unlike Cardinal Bellarmine and Galileo, and Jack and Diane) John and Lise *endorse all the same epistemic norms*, and thus agree entirely on what kind of evidence is relevant to determining whether Signe is guilty of a particular crime (robbery). Lise thinks Signe is guilty. John thinks she is not. If guided entirely by which epistemic principles they accept, they would both reach the conclusion that Signe is innocent. However, in debating the matter with one another, they find themselves disagreeing at nearly every juncture. What accounts for this disagreement? For one thing, each succumbs to a plethora of implicit biases that affect their judgment of the

credibility of one another. For instance, John comes to disproportionately disvalue certain points made by Lise due to his implicit bias that female testimony about stressful matters is overemotional. Lise, for her part, undervalues a range of John's more intellectual points due to her implicit prestige bias, as she knows John lacks a college degree. Their implicit biases, suppose, run deeper and affect many steps. Moreover, the meeting room that's available for them to meet to judge Signe's guilt is at the top of several flights of stairs, causing tiredness that affects judgments; the room also as smells and lighting that trigger other situationally driven cognitive biases which draw them increasingly further apart on the matter of Signe's guilt.

Now, for an observation and a question. *Observation*: John and Lise's disagreement is, in the context described, no less irreconcilable than Jack and Diane's. *Question*: Does this case feature dialogue facts that motivate epistemic relativism? More specifically, does the irreconcilability achieved in the case of John and Lise in any way motivate a move towards epistemic relativism? The answer has to be no. After all, it would be *arbitrary* to think that this case shows that the matter of Signe's guilt should be relativised to John and Lise's respective *epistemic systems*. They are after all, the *same* epistemic system, *ex hypothesi*!

Obviously, what explains John and Lise's inability to resolve their dispute is *not the epistemic systems and epistemic norms they embrace* (which are the same), but rather, the cognitive biases that manifest in their exchanges with one another.<sup>31</sup> The case of John and Lise – even if itself an extreme/atypical case – motivates nonetheless a plausible *constraint* on what kind of dialogue facts are *relevant* to establishing epistemic relativism – *viz.*, that *only dialogue facts that hold in virtue of what epistemic systems one embraces could be relevant to establishing epistemic relativism*.

Now let's revisit the question that was posed earlier: does anything prevent Hales from supposing there are *actual* dialogues that could do the work he's using the Jack and Diane case to do, to support epistemic relativism? The answer, in short, is: only if the kind of irreconcilability (e.g. being such that no Archimedean meta-norm is available) found in the Jack and Diane case can only be accounted for by their endorsing different epistemic systems. After all, this is because: only dialogue facts that hold in virtue of what epistemic systems one embraces could be relevant to establishing epistemic relativism. In the previous section, however, it was shown that empirical evidence of the *extent* and prevalence of cognitive biases gives us good inductive grounds to doubt that, when two individuals (in an actual case) do in fact reach an irreconcilable

position (such that no Archimedean meta-norm is available) this will be a position that is not due, *at least in some part*, to various biases and situational factors which are independent of what epistemic systems are embraced. What follows is that we have good inductive grounds to doubt that there are *actual* dialogues that could do the work Hales is using the Jack and Diane case to do, to support epistemic relativism.

#### 4.8 Retreat to Possibilism?

Let's consider again the Hales-style revamped version of the non-neutrality argument:

11. There can be a non-relative resolution of the dispute between Jack and Diane concerning the soul, only if there is an *Archimedean meta-norm* (e.g. *appropriately neutral* and *appropriately discriminatory*) available.
12. In the context of the dispute between Jack and Diane, no such Archimedean meta-norm is available (ex hypothesi).
13. Therefore, it's not the case that there can be a non-relative resolution of the dispute between Jack and Diane concerning the existence of the moons.
14. Therefore, epistemic relativism is true.

The combined result of the previous two sections is that (12) can be appealed to in support of (14) only against a background commitment to thinking that features of *merely possible dialogues* (e.g. a dialogue where Jack and Diane's disagreement is accounted for *entirely* by their holding different epistemic systems) can underwrite an argument for epistemic relativism. Even shorter: Hales' argument must be understood, *qua* dialogic argument for epistemic relativism, as a 'possibilist' dialogic argument, not an actualist one.

I think there are two central problems with attempting to defend epistemic relativism by appealing to a *possibilist* version of a dialogic argument.<sup>32</sup> The first problem is one of overgeneralisation, and this problem can be highlighted by way of parity. Consider that if one could establish – or more weakly, motivate – epistemic relativism on the basis of properties of possible arguments between possible agents, then there's no good (non-arbitrary) reason to suppose we should not be able to apply the same strategy, *mutatis mutandis*, to marshal support for relativism elsewhere. Take the moral case. We can *imagine* bizarre moral systems, ones which lead imagined interlocutors to a position where there is no

Archimedean meta-norm available. For example, a possible agent, 'Tim', subscribes to a moral system, M1, comprised of moral norms which prohibit certain kinds of discourse, and which mandate other forms of discourse. Obviously, there could be some possible argument between Tim and another possible agent, 'Jim', who subscribes to a fundamentally different moral system, M2, one consisting of moral norms which mandate very different forms of discourse and prohibit other forms. Tim and Jim, suppose, entirely due to their differences in moral systems, are such that no Archimedean meta-norm is available. The point I am after is this: if Hales were to defend (12) on merely *possibilist* grounds, there appears to be no non-arbitrary reason to think that moral relativism couldn't also be supported on the basis of possible disagreements between possible agents, provided those possible disagreements have the same relevant property as the possible disagreement between Jack and Diane – *viz.*, the unavailability of an Archimedean meta-norm.<sup>33</sup> The same point applies to relativism about political systems, science, logic, and the like. Quickly, global relativism begins to raise its head.

We can think of the overgeneralisation strategy as posing a dilemma to the defender of a dialogical argument for epistemic relativism on possibilist grounds.

*Horn 1:* allow for such an overgeneralisation, and accordingly, be prepared to defend a thesis that comes very close to global relativism, a thesis that – in Chapter 2 – was shown to have seemingly intractable problems, *or*

*Horn 2:* explain why the unavailability of an Archimedean meta-norm is supposed to motivate a relativist conclusion when the lack of such a meta-norm is a property of a possible disagreement between Jack and Diane, but *not* when it's a property of a possible disagreement between Tim and Jim.

The first strategy is, if I've made my case in Chapter 2, fundamentally unworkable. The second though looks entirely arbitrary. But a retreat to a possibilist strategy faces a second difficulty. Consider that the class of possible dialogues includes some possible dialogues where the individuals featuring in the possible dialogues are actual, but also, possible dialogues where the individuals featuring in the possible dialogues are merely possible. And then there will be mixed cases: possible dialogues between an actual person (e.g. Bellarmine) and possible person (e.g. 'Galileo').

Plausibly, we'd need to – when claiming that a possible dialogue is one featuring an *actual* person – hold fixed the dialogue-relevant dispositions of the actual person, or else, we are just envisioning a dialogue between two possible people. For example, if we are envisioning a possible argument featuring an actual person such as Bellarmine, we plausibly hold fixed Bellarmine's disposition to (for instance) consult scripture at certain junctures, as this is a dialogue-relevant disposition of Bellarmine. Crucially, though we'll *also* hold fixed the cognitive biases Bellarmine (or any other actual person) has, as these are, as suggested in §4.6, *also* dialogue-relevant dispositions.

But now the issue raised for actualist strategies resurfaces: since irreconcilable positions reached partially due to *cognitive biases* (which are orthogonal to what epistemic systems one embraces) are, as has been argued, not suitable for motivating epistemic relativism, the possibilist strategy would have to make yet a *further* retreat: to possible dialogues where both interlocutors are possible individuals very different from us, individuals who reach their irreconcilable positions *entirely* as a matter of their embracing the epistemic systems they do. But once one retreats from actual to merely possible dialogues with agents very different from us as what's supposed to be doing the work, one (in short) retreats a *very* long way from, say, the attempt to motivate relativism by pointing to actual disputes that proponents of dialogic arguments have traditionally taken to be the relevant ones.

## 4.9 Remaining Problems for the Dialogist

But let's suppose that somehow, a possibilist strategy could be vindicated in the sense that it could be defended that facts about possible dialogues between possible agents very different from us *could* be relevant to whether epistemic relativism is true. There remain two outstanding problems with any attempt to defend a dialogic argument for epistemic relativism. I conclude by discussing them.

### 4.9.1 Resolution, revisited: the gap problem

Let's return now to a fictionalised version of the scene of the crime: early 17th-century Italy. Our characters are 'Balileo' and 'Gellarmino', entirely fictional individuals who embrace the epistemic systems of Galileo and Bellarmine, respectively, but *unlike* Galileo and Bellarmine, are idealised agents entirely free of any cognitive biases, and whose dialectical moves are – unlike ordinary human beings' – *entirely* functions of which

epistemic systems they accept. Balileo and Gellarmino become exasperated with one another, and more generally, with the dialectical impasse that neither seems able to break through, despite adducing as much evidence as either could find.

Suppose Balileo and Gellarmino decide that retreating to scepticism is really just throwing in the towel and decide rather to embrace a *relativist resolution* to their disagreement: P is true relative to Balileo's system, and P is not true relative to Gellarmino's system. Now, as was noted earlier in the chapter in discussing the case of 3 marbles, a relativist *resolution* – that is, a result that entails that epistemic relativism is true – must not *merely* be an agreement to accept conditions such as: *According to Balileo's system, P* and *According to Gellarmino's system, not-P*, as absolutists can accept these conditions. What we said was that a relativist resolution must at least involve an additional element whereby it is denied by both parties to the dispute that there is *any further sense in which each can be right*. Let's now put that extra element in place and suppose that our idealised agents Balileo and Gellarmino, in response to their entrenched, position reason as follows:

**Balileo:** P is true relative to my system, it is not true relative to your system, and there is no further sense in which P is true.

**Gellarmino:** Not-P is true relative to my system, it is not true relative to your system, and there is no further sense in which not-P is true.

Balileo and Gellarmino recognise that their statements are tantamount to epistemic relativism, and so Balileo and Gellarmino, finally, 'resolve' their disagreement by both accepting epistemic relativism – they both believe it is true. And once they do, it does seem they have nothing further to argue about. Accordingly, it's correct to say that *accepting relativism* is as Hales puts it, a 'disagreement elimination' strategy.

But what's doing all the work for epistemic relativism as a disagreement elimination strategy is not epistemic relativism's *being true*, but simply that the two parties to the dispute *believe* epistemic relativism is true. Just as we can disarm our disagreement by believing epistemic relativism is true, we could (for example) disarm our dislike for one another by coming to believe neither of us has wronged the other one. No matter how effective such beliefs are in resolving our dislike for one another, what we've done to wrong one another is insensitive to this. The bottom line posed by the problem, call it the gap problem, is this: epistemic relativism might well be true, but if it is, it's not going to be established by the fact that *believing* it is true can help us to stop arguing; even if all

other hurdles facing dialogic arguments were addressed, a proponent of a dialogic argument – who reasons from the irreconcilability of some possible or actual dispute to the truth of epistemic relativism – looks very much stuck defending the implausible implicit premise that epistemic relativism is established on the grounds that it can be practically useful for some (possible) people to believe it. But it would take a radical neo-Jamesian to try to close the gap between belief and truth on these grounds.

#### 4.9.2 Relativism and scepticism

Even if all other issues raised could be somehow dealt with by dialogic-argument inspired relativist – including the one just discussed – a recurring bogeyman remains: let's suppose for the sake of argument that – in virtue of some dialogue fact (e.g. no available Archimedean meta-norm) – the premise that there can be no non-relative resolution with respect to some disagreement, *D*, is true of *D*. What of such a situation recommends *relativism* over *scepticism*, as the Pyrrhonian originally recommends? That is: what about such situations leaves a Hales-style argument in a *stronger* position than a proponent of a Lammenranta-style dialectical interpretation of the Pyrrhonian argument, construed as an argument aimed to engender *withholding of judgment*, as outlined at the outset of the chapter?

The answer is, not much. And in fact it looks like – at least in some clear respects – the sceptic is poised to claim an important advantage over the relativist. To appreciate this point, it is helpful to juxtapose relativism versus scepticism as a response to clashing dialectical positions with a much-discussed and more refined version of this problem as pursued by contemporary social epistemologists. The matter of what the reasonable response is in the face of disagreement is, along with debates about testimony and transmission, perhaps the most hotly debated contemporary issue in social epistemology.<sup>34</sup> And so it will be instructive to consider the contemporary formulation of the kernel of the philosophical problem: is doxastic revision rationally required in the face of a recognised *peer* disagreement?

The fact that this contemporary debate in social epistemology is framed in terms of *peer* disagreement, rather than disagreement simpliciter, has a simple explanation: consider that if you disagree with someone obviously less well informed than you or otherwise less likely to be right on the matter than you, it's plausible to think that a perfectly reasonable response is to 'hold your guns'. Things become much more interesting when the disagreement in question is between individuals regarded as

epistemic peers, on a par in terms of their evidence and abilities as these bear on the target question. When disagreements have this feature, the received way of carving up the space of options in the literature offers a choice between the *conciliatory* view, which says that in such circumstances, it's rationally appropriate to lower our credence in the target proposition – or as Feldman has suggested, to withhold judgment; and, the *steadfast* view, on which it can be appropriate to hold one's guns in the face of a recognised peer disagreement.

Two points are worth noting here. Firstly, although on the steadfast view (e.g. Kelly (2005)) each party *x* is rationally permitted to walk away thinking *x* is correct (respectively) – something that *seems* to have some of the trappings of an 'everyone wins' scenario often associated with epistemic relativism – the disagreement-departure scenario that follows from the steadfast view is a very different one indeed. Note, importantly, that on the steadfast view, it's going to follow that at least one party has a false belief that he/she is rationally permitted to hold. Nothing recommended by the steadfast view suggests that there should be, or even that either party should believe there should be, no further sense in which either party is right beyond just right according to each perspective.

According to the contrasting position, *conciliatorism*,<sup>35</sup> it is *not* rationally permissible to hold one's guns in the face of a recognised peer disagreement. One must (for instance) 'split the difference' between one's credence and one's interlocutor's credence in the target proposition (e.g. Elga 2007) or simply withhold judgment – and accordingly embrace agnosticism – with respect to the originally believed proposition (e.g. Feldman 2007).

The conciliatory view is not, in and of itself, a 'dialogic argument for scepticism' (one that moves from facts about particular dialogues to the conclusion that scepticism is true). *However*, one of the stock objections to the conciliatory views is that it has, as a consequence, a rather wide-sweeping kind of scepticism. As Clayton Littlejohn (2013, 171) has put the worry:

Most of the interesting things we believe (i.e., most of what we believe about epistemology, ethics, metaphysics, politics, and religion) are controversial. Much of this controversy seems to involve peers who disagree with each other fully aware of the fact that there are peers that they disagree with. Because [conciliatorism] is correct, we cannot rationally remain committed to these controversial propositions. Thus, few of the interesting things we believe we believe

rationally: ...The pessimistic conclusion is that we should suspend judgment on most of the interesting things we believe.

I do not want to claim here that conciliatorism is the right response to peer disagreement, or even for that matter that *if* conciliatorism is true, then the kind of scepticism Littlejohn intimates is unavoidable. Rather, the point that should be highlighted is that there is, *a la* Littlejohn, an entirely straightforward pathway from dialogic facts to the conclusion that at least a version of scepticism is true. Scepticism is a position after all about the scope of human knowledge; radical scepticism maintains there is no human knowledge. Less radical views claim, to different extents, that there is less knowledge than we ordinarily take for granted.

While the steadfast view speaks to neither relativism nor scepticism, the conciliatory view, if correct, highlights a way that dialogic facts could potentially 'close the gap' (or, at least, narrow the gap) which seemed intraversable for the relativist. The 'gap', recall, was that, *even if* as a consequence of how a dialogue goes, two parties eliminate disagreement by *believing* that epistemic relativism is true (a strategy that would, granted, be effective in *eliminating* disagreement), this much – absent a collateral fact constructivism in hand – has no import for whether epistemic relativism actually *is* true. But the same is not so in the case of scepticism, in the following sense: *if*, as a consequence of how a dialogue goes, rationality requires that both parties withhold judgment on the target proposition (as the conciliatorist would have it), it follows that neither walks away with knowledge, provided rational belief is required for knowledge. Putting this all together: not only is there not a clear reason to recommend relativism *over* scepticism as a response to certain kinds of deeply entrenched arguments, but moreover, there seem to be positive reasons for accepting that the sceptic has a much more coherent story to tell about how the sceptical position could be motivated by dialogic facts than the relativist has.

#### 4.10 Concluding Remarks

Although dialogic arguments are among the most popular argument-type for epistemic relativism, they aren't – as I've shown here – ultimately promising for motivating epistemic relativism. I began by considering the most simple version of such an argument, the 'no-neutrality, therefore relativism' argument, which attempted to move from a certain property of an actual dialogue between Galileo and Bellarmine – namely, that

there seemed to be no appropriately neutral meta-standard to which either could appeal to resolve their dispute – to the intermediate conclusion that there can be no non-relative resolution to the disagreement – a claim that at least appears to bring one closer to epistemic relativism. According to Siegel's characterisation of this argument strategy on behalf of the relativist, the 'non-neutrality' doing the work is regarded as a consequence of there being (in cases such as Galileo v. Bellarmine) both a first-order disagreement *and* a second-order disagreement about what kinds of evidential standards are relevant for adjudicating the first-order issue.

I highlighted an easily overlooked point which requires a 'tweak' to any version of a 'no-neutrality, therefore relativism' argument: consider that although there surely are *sufficiently neutral* meta-norms either could appeal to (e.g. consider the norm: *if A, infer A*), sufficiently neutral meta-norms might not also be sufficiently *discriminatory* for the purposes of settling the dispute. The 'no-neutrality, therefore relativism' argument, suitably sensitive to this point, should be framed in terms of a lack of what I called an *Archimedean* meta-norm (a meta-norm that is appropriately neutral *and* appropriately discriminatory), and thus: the best way to understand this simple argument should be as relying on the premise that: in virtue of Galileo and Bellarmine disagreeing at both the first-, *and* the second-order, there could be no Archimedean meta-norm (and, not merely, no appropriately *neutral* meta-norm) available to them.

My next move in this chapter was to show, by parity of reasoning, that if Galileo and Bellarmine really *were* in such a position that there was no Archimedean meta-norm available to them, this *wouldn't be* (as the simple version of the argument has it) *simply* because Galileo and Bellarmine disagreed at the first- and at the second-order. The parity argument drew attention to a contemporary debate about knowledge-how, wherein we have first- and second- order disagreement, but where there was *no* good reason to think no Archimedean meta-norm was available. The upshot was that if the famous Bellarmine/Galileo clash was such that there was no Archimedean meta-norm available, this must be in virtue of something *in addition* to the mere first- and second-order disagreement highlighted by the simple version of the argument.

Of course, there is a way around this issue: we might advert to a case where the unavailability of an Archimedean meta-norm is *built in to the case* – we find such a case envisioned in recent work by Hales (2014). In Hales' case, whatever conditions over and above first- and second- order

disagreement that would be needed to secure the unavailability of an Archimedean meta-norm are *de facto* satisfied, as he stipulates the interlocutors' position is irreconcilable (a position that *wouldn't* be reached if an Archimedean meta-norm *were* available). Hales' version of the case constitutes probably the strongest showing for a dialogic argument. Though, as I argued in the remainder of the chapter, intractable problems lay waiting.

The structure of the remainder of the chapter can be summarised very simply. I argued that even the most promising looking dialogic arguments (such as Hales' revamped version of the 'no neutrality, therefore relativism argument') must be defended *either* on actualist or possibilist grounds – that is, by adverting to properties of actual or merely possible dialogues. I then argued – over the course of the remainder of the chapter – that neither strategy is workable, and concluded by suggesting that even if the problems I'd raised for actualist and possibilist strategies, as such, could be overcome, there remained two further problems – the first was what I called the 'gap problem' and the second was a recurring problem that also faced the rationale Sankey puts forward on behalf of the epistemic relativist: namely that, at the end of the day, there is going to be no compelling reason to embrace relativism rather than scepticism.

I want to conclude by quickly addressing what might seem like a glaring oversight: how can I dismiss dialogic arguments for epistemic relativism without even engaging with *faultless disagreement-style* arguments for epistemic relativism, of the sort defended by Kölbel (2004) and MacFarlane (2007)? The first part of the answer is: don't worry, I'll be engaging with semantically driven arguments for epistemic relativism in a later chapter (i.e., Chapter 7). But the more substantive answer is: faultless disagreement-style arguments are *not* dialogic arguments for epistemic relativism, despite the many superficial similarities. MacFarlane and Kölbel-style faultless-disagreement arguments simply do not regard properties of any *particular* disagreement as even *the sort of thing* that's in the market for establishing epistemic relativism. Though MacFarlane and Kölbel could indeed point to the same kinds of cases as Hales does in making the point they're after, what MacFarlane and Kölbel are actually doing is reasoning from semantic and pragmatic evidence about disagreement *patterns*, much more generally, to the conclusion the a relativist semantics (in certain domains where we find such disagreements) best explains our practices of attributing certain terms. One aim of this book will be to ultimately show that semantic arguments for epistemic relativism, of the sort of which faultless-disagreement style arguments

are an instance, do not fall prey to some of the familiar patterns of problems betrayed by more traditional (non-semantically motivated) arguments for epistemic relativism. But before we engage with such semantic arguments, there is another more traditional argument strategy that needs to be engaged with, one which adverts in the main to incommensurability and circularity. This will be the topic of Chapter 5.

# 5

## Incommensurability, Circularity and Epistemic Relativism

*Abstract.* This chapter outlines and evaluates a strategy for motivating epistemic relativism which draws from considerations to do with the incommensurability of epistemic systems. Such arguments can usefully be understood as beginning in the same place as dialogic arguments. But rather than to attempt, as dialogic arguments did, to establish the epistemic relativist's conclusion by appealing to dialogue facts, incommensurability arguments attempt to get there *via* appeal to epistemic circularity. Once the structure of these arguments is suitably sharpened, it is shown that the variety of circularity such arguments betray bears some tight commonalities with a variety of circularity that is most often discussed in connection with contemporary debates about perceptual warrant (e.g. Pryor 2000; 2004; Wright 2008). It is concluded that, while incommensurability arguments can lay claim to a serious threat to the *cogency* of our attempts to justify our epistemic principles, this threat to cogency fails to motivate epistemic relativism over dogmatist, conservative or sceptical alternatives.

### 5.1 Back to the original position

In the previous chapter, we saw that *dialogic arguments* for epistemic relativism did not look very promising. The passage to epistemic relativism that this approach opted for, one that highlights properties of actual or possible dialogues and tries to move from these properties to epistemic relativism, seemed intractably blocked, and at multiple places.

Maybe though, dialogic arguments started out in the right spot – for the purposes of motivating epistemic relativism – but simply took a wrong

turn from there. For ease of exposition, let's use 'the original position' to refer to a gambit-point that is familiar from the previous chapter: the position where A finds herself attempting (unsuccessfully) to settle the matter of  $p$  with B, and furthermore, where no Archimedean meta-norm looks to be available either to appeal to or to rationally break through the deadlock.<sup>1</sup> This is effectively the position that Hales (2014) tells us Jack and Diane are in, when Jack consults analytic philosophy of mind and Diane consults the Catechism, when attempting to establish the nature of the human soul.

The bane of dialogic arguments was, as we saw, that – at least, in virtue of what I called 'dialogue facts' – no actual or possible dialogues plausibly motivate epistemic relativism. This was the thesis of Chapter 4 at any rate. But perhaps, from the original position, there is more effective move in the direction of epistemic relativism, one which does not for its relativism-motivating effectiveness depend on how any actual or possible dialogues go. The sort of argument we'll now explore proceeds in this fashion.

The topic of this chapter is *incommensurability arguments* for epistemic relativism. While incommensurability has been used in a number of ways, the core datum adverted to is (in very broad terms) *no common measure*. Although the term originated in Ancient Greece to refer, specially, to the mathematical notion of no common measure between magnitudes,<sup>2</sup> arguments under the heading of incommensurability have highlighted various different senses in which it seems as though (broadly speaking) from two perspectives, there can be no common measure, be it semantic, conceptual, or in the cases we're especially interested in here – epistemic.<sup>3</sup>

Among others, Wittgenstein (1969), Kuhn (1962) and Hacking (1982) have proposed views on which – for different reasons, in each case – individuals can find themselves in a position where it looks as though neither party is in a position to rationally persuade the other, given an apparent lack of any discoverable Archimedean meta-norm (in the sense articulated in Chapter 4). And from *here*, incommensurability arguments for epistemic relativism attempt to proceed through a passage that bypasses *dialogue* facts – that is, incommensurability arguments use the original position as a kind of 'prop' for motivating what looks like an inescapable circularity problem which *itself* (rather than any facts about how one interacts with one's interlocutor) should be understood as what's doing the heavy lifting in the service of generating the result that epistemic justification is at most system-relative.

Pared down more simply: the boilerplate difference between the two strategies is that what's primarily relevant to motivating epistemic relativism, for the kind of dialogic arguments surveyed in Chapter 4, is that we *and our interlocutor* cannot break out of a deadlock in which we find ourselves in the original position; that is, it's a position the holds because of *relational properties between interlocutors*. By contrast, what's of primary relevance to motivating epistemic relativism, for incommensurability arguments is that once we're in the original position – a position from which it is incumbent upon us to attempt to justify our own epistemic system – we are left with nothing to show for ourselves but a kind of inevitable 'circular' justification which it seems we cannot break out of.<sup>4</sup> The incommensurability–circularity–relativism sequence thus takes the following shape:

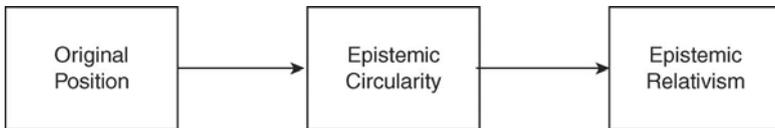


Figure 5.1 Incommensurability Arguments

Here's the plan. This chapter will explore the viability of this kind of an argument-type for epistemic relativism by exploring, in order, three central questions about the incommensurability–circularity–relativism sequence:

**Q1:** How does epistemic circularity *arise* from epistemic incommensurability?

**Q2:** What is the nature of this circularity?

**Q3:** How is this kind of circularity supposed to get us to epistemic relativism?

## 5.2 How does epistemic circularity arise from epistemic incommensurability?

It's not an accident that we aren't often in a position – i.e. 'the original position' – where we actively try to provide justification for the epistemic principles we rely on, such as inference, deduction, observation, etc. This is because we ordinarily share a common background with our interlocutors, where the differences that are the source of typical

epistemic disagreements concern the application of shared epistemic principles in particular cases, and *not* whether these very principles are themselves correct. In short: we don't always agree about what inference and deduction *point* to, in particular cases, but preaching about the epistemic merits of deduction and inference is almost *always* preaching to the choir.

When we are in the original position, things may well be relevantly different. When, for instance, Bellarmine and Galileo find themselves so diametrically opposed, *vis-à-vis* *p*, such that they can't seem to agree about what basic kinds of evidence are relevant to adjudicating whether *p*, the prospect of trying to justify our own epistemic systems – unlike in ordinary circumstances, where we can take a shared background for granted – thereby gains a kind of *relevance* that such an activity would otherwise seem to lack entirely.

In outlining how epistemic circularity arises from incommensurability – *viz.*, that is, in developing an answer to Q1 – I want to first examine the following: how incommensurability generates a special kind of enabling condition for epistemic circularity to arise: by being the sort of communication context where there is a certain *relevance* (both practical, and as I'll also suggest, epistemic) to justifying principles which are ordinarily pointless to try to justify.

Consider first this point in terms of purely *practical* relevance, setting epistemic relevance aside for now (we'll return to this). Transposed to terms familiar from Stalnaker (1978): original-position scenarios, in virtue of bringing into contact with one other incommensurable epistemic frameworks, give rise to a unique kind of communication context where the set of 'live options' (in these communication contexts) includes the *denials* of epistemic principles that each interlocutor ordinarily takes for granted as in the common ground.

While attempting to establish the truth of one's own epistemic principles is, in nearly every single communication context we face, 'pointless' in light of our communicative objectives in that doing so *de facto* does not cut down any *live options* – the same is simply not so in original-position-style scenarios. Original-position-style scenarios thus, in virtue of what propositions remain *live* (in light of the opposing frameworks being taken for granted by the interlocutors) furnish a special kind of practical relevance to justifying epistemic principles where such practical relevance seems on reflection *distinctive* of cases like what we find at the original position.

And of course, if it turns out that the justifying of epistemic principles that appears practically relevant in such scenarios leads to a circularity

problem that *itself* marshals support for epistemic relativism, then we can see how the path articulated in Figure 5.1 might work: a path that *begins* at the original position – and which moves, from this position, through circularity rather than dialogue facts – to epistemic relativism.

But one might balk at the shape of this move, so described. While the Stalnakerian line speaks to the practical relevance of attempting to justify something we wouldn't ordinarily seek to justify – our own epistemic principles – it might seem that a more viable passage from incommensurable frameworks, to circularity, to epistemic relativism, will have to identify at least some way in which the confronting of radically different epistemic frameworks might bring about not merely a practical, but an *epistemic*, relevance, to justifying our own epistemic principles.

Let's now consider how such a line might well proceed. Here it will be useful to briefly consider some of considerations that have motivated various genealogical arguments for scepticism about certain kinds of philosophical judgments. In abstract form, genealogical arguments first highlight some aspect of the genealogical contingency of one's (generally, evaluative) beliefs and on the basis of this observation challenge the epistemic justification of the beliefs in question. One familiar way to try to do this is to attempt to show that the causal origins of one's beliefs constitute an epistemic defeater for the belief in question, as is the standard move advanced by evolutionary debunking arguments (e.g. Joyce (2007)).<sup>5</sup>

Interestingly, though, another strand of genealogical argument has been brought to the fore by experimental philosophy, which, as Amia Srinivasan (2014, 3) observes, is:

largely devoted to arguing that people's judgments about epistemology, ethics, philosophy of language and metaphysics systematically vary with culture, gender, socioeconomic status and extent of philosophical training.

Suitable appreciation of differences, including radical differences, in epistemic perspectives – an appreciation we can't avoid when in the original position – might well, in certain circumstances<sup>6</sup> cause what Knobe and Nichols (2008, 11) describe as a:

[...] crisis akin to that of the [Christian] child confronted with religious diversity... For the discovery of religious diversity can prompt the thought that it's in some sense accidental that one happens to

be raised in a Christian household rather than a Hindu household. This kind of arbitrariness can make the child wonder whether there's any reason to think that his religious beliefs are more likely to be right than those of the Hindu child.... And just as some Christian children come to think that there's no rational basis for preferring Christian to Hindu beliefs, we too might come to think that there's no rational basis for preferring Western philosophical notions to Eastern ones.<sup>7</sup>

Resituating the Knobe–Nichols point in terms of *epistemic* perspectives (as per the original position): the very *confrontation* of radically different epistemic perspectives or systems might render the provision of justification for our epistemic systems not merely practically relevant but *epistemically* relevant with respect to our own epistemic systems; and this for the reason that such confrontation can potentially (and not irrationally) prompt the thought that it might be in some way accidental that one has come to embrace the particular epistemic system that one does.<sup>8</sup>

But of course some individuals will remain entirely unmoved by such considerations and will accordingly not – as Knobe and Nichols' envisaged Christian child confronted with religious diversity does – form as a result of such an encounter a *doubt* that could itself function as a psychological defeater (which is then incumbent on the individual to defeat). But even for such individuals, there remains a rationale for supposing the very confrontation, in such a context, might well itself be epistemically significant in the sense that it makes epistemically relevant one's justifying one's own epistemic principles.

In order to unpack this point, compare the epistemic significance of such confrontation with the epistemic significance of becoming *aware* of an error possibility. One of the most interesting observations about relevant alternatives theories of perceptual knowledge (e.g. Dretske (1970)) concerns the matter of what kinds of error possibilities are epistemically *relevant* in the sense that knowledge of some target proposition requires the ruling out of such possibilities that are incompatible with it.

As Pritchard (2010, 260–263) has noted, any plausible version of a relevant alternatives account of knowledge must recognise (at least) two entirely different ways an alternative can qualify as epistemically relevant. One way is by obtaining in near-by worlds, *regardless* of whether one is aware of the error-possibility in question. This will be as things stand in epistemically inhospitable environments, such as the infamous

'barn façade county'. Though another very different way an alternative can qualify as epistemically relevant is by one's simply being *made aware* of the error possibility and its incompatibility with the target proposition (e.g., Lewis 1996).

For instance, if one is at a zoo, staring at zebra-looking creatures, and then (perhaps, due to overhearing a conversation about potential pranks and the mule-replacing plans they've devised) becomes aware of the 'cleverly disguised mule hypothesis' and its incompatibility with the target proposition (i.e., that one is looking at a zebra), it seems one must thereby be able to *rationaly dismiss* the cleverly-disguised mule hypothesis in order to retain her knowledge that she's looking at a zebra. And, at least as Pritchard's story goes, one might rationally dismiss such an alternative *either* by *discriminatory* or by *favouring* epistemic support, where the latter kind of support can include such things as background evidence rather than discriminatory abilities.

Putting this all together, an answer to Q1 begins to unfold: (i) *ordinarily*, given what is taken for granted by one's interlocutors, it is practically pointless to attempt to provide any rational support for our own epistemic systems; however, the unique starting point that features in the original position – where radically divergent epistemic systems are brought into contact with one another – is an exception in that (ii) there is a *practical* relevance to defending one's epistemic system that is unique to this kind of situation for the reason that the communication context that corresponds with the original position is one where the *denials* of our own epistemic systems are live possibilities. But even *more* importantly (iii) there is an *epistemic* relevance to doing so, where the epistemic relevance can be defended in either of two ways. Firstly, (iii-a), as per the Knobe–Nichols point: given that radically different epistemic systems will typically correspond with radically different cultural backgrounds, the encountering of radically divergent epistemic systems can (at least potentially) engender doubt about whether there is in fact a rational basis for preferring (for instance) our Western epistemological principles to radically different systems employed by individuals who happen to have ossified their epistemic policies through very different kinds of practices. And such doubts can underwrite the epistemic relevance of attempting to justify one's own epistemic principles.<sup>9</sup> Secondly, (iii-b), the point about epistemic relevance can be sharpened by drawing the analogy to the relevant alternative literature – where there is a straightforward precedent for regarding mere *awareness* of an incompatible error possibility as one of two ways (along with the obtaining of that possibility in near-by worlds) that suffices to make that error possibility

epistemically *relevant* – such that, awareness of the error possibility requires one to be able to rationally dismiss the error possibility in order to retain one’s knowledge of the target proposition.<sup>10</sup>

The bit sketched so far captures the ‘enabling condition’ aspect of Q1: we have a reason why something very unusual – the activity of justifying our own epistemic principles – might happen very naturally off the back of confrontations with radically different epistemic principles. But to get from incommensurability to circularity, we need see how *circularity* is supposed to appear, once the activity of justifying one’s epistemic principles gets off the ground.

Probably the most succinct statement on this score has been offered Michael Williams (2007, 3–4), who paints the following picture of how things go *once one attempts to justify* one’s own epistemic framework:

In determining whether a belief – any belief – is justified, we always rely, implicitly or explicitly, on an epistemic framework: some standards or procedures that separate justified from unjustified convictions. But what about the claims embodied in the framework itself: are they justified? In answering this question, we inevitably apply our own epistemic framework. So, assuming that our framework is coherent and does not undermine itself, the best we can hope for is a justification that is epistemically circular, employing our epistemic framework in support of itself. Since this procedure can be followed by anyone, whatever his epistemic framework, all such frameworks, provided they are coherent, are equally defensible (or indefensible).

Obviously, Williams is gesturing here, albeit, in broad-brush strokes, toward how the incommensurability–circularity–relativism sequence unfolds. But let’s focus on how his remarks help address Q1: how circularity, specifically, is meant to *arise* within the incommensurability–circularity–relativism sequence. Williams’ articulation of the relevant strand of reasoning here (as it bears on Q1) is as follows: Once one (S) is in a position where S is trying to justify S’s own epistemic framework or system, X, by attempting to justify the claims that comprise the system ( $x_1 \dots x_n$ ), then:

- (i) *Inevitable application*: S must (inevitably) *apply* that system (X).
- (ii) *Application-circularity*: the application, by S, of a system X to justify the claims ( $x_1 \dots x_n$ ) of that very system, X, is sufficient for leaving S’s epistemic justification for the claims of X ( $x_1 \dots x_n$ ) circular.<sup>11</sup>

These abstract claims, *inevitable application* and *application-circularity*, lie at the heart of several well-known epistemic incommensurability arguments which have been used to motivate epistemic relativism.

For Wittgenstein (1969), in *On Certainty*, we can think of inevitable application and application circularity as arising *via* the potential incommensurability of different ‘hinges’, or bedrock foundational propositions. On Wittgenstein’s brand of epistemology – outlined in Chapter 2 – any kind of rational assessment of beliefs takes place against a background of a commitment to hinges that *must* stay put (e.g. it is part of the logic of our investigations for Wittgenstein they lie outside the scope of rational doubt, from within the system) in order to make possible any kind of rational assessment.<sup>12</sup> So while Wittgenstein’s particular brand of epistemology can’t make room for rationally doubting our own hinges, we can easily imagine one (in an original-position-type situation) drawing attention to *characteristics* of their own hinges which are being put forward as in some way ‘epistemically better’ than one’s interlocutor’s hinges.<sup>13</sup> Quickly, though, we get to *inevitable application*: Wittgenstein regarded our trust in our hinges, at the bottom of our language games, as a ‘form of life.’<sup>14</sup> Attempting to suggest one’s ‘form of life’, or the assessments implied by it, are epistemically *better* than another’s will inevitably be made from *within* the *standards of evaluation* that are distinctive of one’s particular form of life<sup>15</sup> – and thus, application circularity.<sup>16</sup>

Granted, the matter of whether individuals can actually, in practice, differ radically in their hinges is one that is disputed among Wittgenstenian commentators.<sup>17</sup> Interestingly, Ian Hacking’s (1982) path to epistemic relativism *via* incommensurable ‘styles of reasoning’ manages to bypass entirely this point. Hacking’s key line, in summary, is that while some statements can be made in any language, others require what he calls a *style of reasoning* such that what ‘is a candidate for being true-or-false’ depends on whether we have ways to reason about it, because ‘what is true-or-false in one way of talking may not make much sense in another until one has learned how to reason in a new way’ (Hacking 1982, 331).<sup>18</sup> Accordingly, for example, ‘statistical reasons had no force for the Greeks’ much like reasons offered in some ancient systems are incomprehensible, as reasons, *vis-à-vis* contemporary science.<sup>19</sup> Examples of such ways of reasoning include ‘renaissance medical, alchemical and astrological doctrines of resemblance and similitude [which] are well-nigh incomprehensible...the way propositions are proposed and defended is entirely alien to us’<sup>20</sup> (1982: 330). On Hacking’s line, in a circumstance where a contemporary Western

scientist and a renaissance alchemist confront one another, and are accordingly led to attempt to justify their respective epistemic systems, each will be relegated to appealing to reasons which have a *sense* only within the style of reasoning that is licensed by the system they already embrace; and, accordingly, any attempt to *justify* one style of reasoning over another will (inevitably) be made *via* appeal to reasons that are given their sense by the very style of reasoning they are attempting to justify<sup>21</sup>.

A parallel situation unfolds when we run a similar sort of ‘original-position’ case, framed in terms of Kuhnian paradigms – where a scientific paradigm is to be understood as including a particular epistemic framework (which will include epistemic principles) along with distinctive concepts and methods characteristically employed within that framework.<sup>22</sup> In the Kuhnian version of an original-position-style case,<sup>23</sup> where one finds oneself in the rare position of it being neither practically nor epistemically irrelevant for one to attempt to provide epistemic justification distinctive of the paradigm within which one is operating, one’s doing so will inevitably be from *within* one’s own epistemic framework (employing the distinctive concepts and methods of one’s own framework); and from this inevitable application we get *application circularity*: one is, in engaging in such a task, appealing to one’s own framework in order to justify the claims that constitute it.<sup>24</sup>

### 5.3 What is the nature of the circularity in question?

So in the face of confrontation with principles, frameworks, styles of reasoning, hinges, etc., from which there seems to be little to no ‘common measure’ with our own – *viz.*, where it is not *irrelevant* (practically, or epistemically, in the sense of §5.2) to justify our own epistemic principles, we’ve seen at least in rough terms how circularity is supposed to arise; it seems that our activity of justifying will be, inevitably, a ‘self-congratulatory pat on the back’.<sup>25</sup>

But epistemic circularity is a genus with multiple species. And some versions of it are more objectionable than others. In this section, the aim will be to clarify the *nature* of the epistemic circularity one faces when attempting to justify one’s own epistemic principles by (in various ways) appealing to them.

We’ll proceed in two steps. The first step will be to consider whether we’ve already been too quick. Might circularity be avoided altogether by simply ‘going externalist’– *viz.*, by, and in a way broadly analogous to what reliabilists are committed to in the case of perceptual knowledge,

supposing that our beliefs *about* our own epistemic principles might be justified *so long as* our beliefs about them are reliably produced?

After engaging with this point, the remainder of this section develops an answer to Q2 by sharpening and refining the nature of the kind of epistemic circularity that should be regarded as ‘in play’ in the incommensurability–circularity–relativism sequence.

### 5.3.1 Interlude: externalism, circularity and bootstrapping

One might be inclined to try to nip things in the bud by ‘calling foul’ and insisting that what’s been described thus far looks like a rigged game – one on which the arising of circularity can be contested as ‘theory-laden’, in a way that simply begs the question against the externalist. Why not think, after all, that the claims that comprise our very epistemic framework – *viz.*, our epistemic principles – are justified *provided they satisfy some externalist condition*, such as reliability?

Consider, for instance, an analogy to the perceptual case. Standard (flat-footed) process reliabilism maintains that our perceptual beliefs can be known so long as they’re the product of a reliable process. And this is so even if we don’t have some additional piece of knowledge (or justification) that the process *itself* is reliable. Call this commitment associated with reliabilism *basic knowledge*:

*Basic Knowledge* (BK): S knows that *p* even though S has no antecedent knowledge that the process R that produced S’s belief is reliable.<sup>26</sup>

Consider now why a proponent of BK looks as though she can neatly sidestep an obvious sort of (analogous) circularity problem: for if BK is *false*, and one needed to, say, justify the reliability of one’s belief-forming process in order to attain justified beliefs *vis-à-vis* the deliverances of the process, then it looks like one falls into the ‘self-congratulatory’ trap: one applies a method to justify that method (e.g., as were if one attempted to demonstrate *that* perception was reliable by using perception to form beliefs). And this looks very much akin to using a legal text to prove that that very legal text is lawful.

But reliabilism, by embracing BK, looks to avoid such a trap altogether. And this is because the kind of externalist justification reliabilists allow for the deliverances of perception is already in place independent of any further facts about the character of the reasons which could be adduced to support the reliability of perception. This is relevant because we can now envision a parallel move at the level of epistemic *principles*.

Let's substitute for (*p*) the claim that one's own epistemic principle (EP) is correct. Call the resulting view Basic Knowledge-Epistemic Principles (BKEP):

*Basic Knowledge-Epistemic-Principles* (BKEP): S knows that epistemic principle EP, which S accepts, is correct, even though S has no antecedent knowledge that the process R that produced S's belief that EP is correct is reliable.

BKEP is in fact just an instance of BK. Of course, some proponents of BK will want to restrict the class of propositions that can be substituted for (*p*) in BK. However, we can easily imagine BKEP as attractive for the wholesale reliabilist. The interesting thing about BKEP is that it looks like an externalist, by embracing it, can lay claim to immunity from the kind of 'application circularity' issue in precisely the same way a proponent of BK can when vindicating (for example) basic perceptual knowledge. After all, if what gives *any* belief positive epistemic status is properties of the *source* of the belief (e.g. its reliability) then – if reliabilism is true – beliefs *about* our epistemic principles can enjoy positive epistemic status provided their source of production is reliable.<sup>27</sup> I have three remarks to make about such a line before moving on.

Firstly, an appeal to BKEP is an instance of an *undercutting* strategy; epistemic circularity is not regarded as arising, and then *overcome* (this would be an overriding or rebutting strategy); rather, it's regarded *not arising* in the first place. The problem is, an externalist/reliabilist reliance on BKEP might well beat application circularity down with a mallet, only to have another variety pop up elsewhere.

The other variety is an artefact of *bootstrapping*, an arguably illicit<sup>28</sup> line of reasoning where (on the assumption that reliabilism is correct) one can – to put the idea in terms of perception – acquire perceptually grounded track-record evidence that one's own perceptual faculties are reliable.<sup>29</sup>

This specific variety of circularity that materialises off the back of bootstrapping (*via* BKEP) would proceed as follows. With reference to BKEP, S can know that epistemic principle EP, which S accepts, is correct, even though S has no antecedent knowledge that the process R that produced S's belief is reliable. For simplicity, let's substitute for 'EP' a simple epistemic principle Boghossian (2006a, 64) calls *Observation* (O): For any observational proposition *p*, if it visually seems to S that *p* and circumstantial conditions D obtain, then S is *prima facie* justified in believing *p*.

Now, the reliabilist who adverts to BKEP to say that (O) can be known, *even in the absence* of any antecedent knowledge/justification that the

process that produced (O) is reliable, must grant that (provided S's introspection is reliable) S can *also* know that it was a combination of (at least) induction and perception along with observation (call this combination O+I+P) which produced S's true belief that (O). But then, awkwardly, S can then infer deductively that S's O+I+P has on this occasion generated a true belief. Repeat this process and attain track-record evidence that, on a number of occasions O+I+P are reliable. But *this* looks illicit to reach, for the same reason it seems illicit to rely on perception to conclude by track record evidence that perception is reliable. It is, at any rate, a *form* of epistemic circularity. The first point thus is that we should be suspicious of the externalist's advertisement as one of a thorough-going *undercutting* strategy.

Of course, the proponent of BKEP might well say that even if (*via* BKEP) a kind of circularity pops up elsewhere, this emergence of such circularity is *compatible* with upholding the key claim doing the undercutting work: and this is the claim that justification for one's epistemic principles can persist in the absence of any application of one's epistemic principles in the service of justifying them. BKEP entails there need be no 'inside job.'

The second point I want to make is that we should be at least somewhat suspicious of *this* claim also. And this is for two related reasons. Firstly, there is an obvious respect in which this is a self-congratulatory pat on the back, albeit in the case of reliabilism, one that surfaces for the *theorist* even if not for the agent, in practice. Whereas an implication of BKEP is that S's belief that her own epistemic principle can be known/justified in the absence of S appealing to any principle to justify this, the *theorist* who is committed to BKEP, in her account of what makes S's belief justified, makes explicit the connection between the principle being justified and the reliable process in virtue of which the agent's belief about her own epistemic principle, is justified. Application circularity can materialise here even when there is no explicit application at the object level.

More interestingly, though – and this is the second reason why we should be suspicious of the reliabilist's adverting to BKEP to avoid an 'inside job' charge – it *might* well be that the undercutting strategy (suitably understood) is not as simple as it appears. One idea that should be dispelled immediately is this: even if, as per reliabilism, one is justified – and indeed, even if one counts as knowing – that one's epistemic principle is correct *provided* the processes leading to the formation of the belief that the principle is correct (whether they be reasoning in accordance with epistemic principles or not) are reliable, it doesn't follow that

one is thereby epistemically positioned to *assert* that one's epistemic principle is correct. And as Mikkel Gerken (2012, 379) has suggested in recent work, while *some* conversational contexts are ones where 'S may assert something although S is unable to provide any reason for it', what he calls *discursive contexts* are not permissive in this way.<sup>30</sup>

Gerken terms *discursive conversational contexts* as ones where 'interlocutors share a presupposition that an asserter must be able to back up unqualified assertions by reasons...' and in which 'being a cooperative speaker involves being sensitive to reasons for and against what is asserted' (2012, 379). In such contexts, *dogmatic* assertions violate a plausible constraint on epistemically appropriate assertion: namely, that such an assertion be *discursively justified*, where discursive justification is something S possesses only if she is able to articulate some epistemic reasons for believing that p. The very *features* of the original position which I suggested earlier render the provision of reasons for one's own epistemic principles uniquely (in this context) practical and epistemically relevant (in the sense of §5.2) are indicative of a discursive context, a context where dogmatic assertion is impermissible. To the extent that the original position really does make practically and epistemically relevant the activity of justifying, we shouldn't assume – and this is the misconception worth dispelling – that dogmatic assertion of one's epistemic principle (in such a context) simply renders *unnecessary* any provision of further reasons for one's principle. On a rationale like Gerken's, the more accurate way to put things might well be that if the warrant one has for one's epistemic principle lies beyond what one can articulate, then the following is true: one simply isn't very well positioned, epistemically, to do something (justify) that seems both practically and epistemically relevant to do. In this way, we should caution against thinking that one can simply *rely* on BKEP to explain why we shouldn't regard circularity as arising from the original position in the first place.

The third point that I think is worth highlighting is that *even if* one could overcome the variety of circularity that arises off the back of bootstrapping, it would not follow that one will have *thereby blocked* the variety of circularity that arises in the incommensurability–circularity–relativism sequence.

Consider the following example, an imagined twist on Evans-Pritchard's (1937) famous discussion of the Azande:

ZAZANDE: In a more remote area of the Congo than where the Azande<sup>31</sup> reside, we find the Zazande. Whereas the Azande valued the Poison Oracle in part because the Azande believed *that the poison*

*oracle was reliable*, the Zazande afford reliability no such importance. The Zazande, given their unique motivational structure and distinctive set of desires, seek only to find what they call in their culture *The One Great Truth* and are uninterested in the accuracy of trivial matters (which include, given their distinctive desires, everything except for the One Great Truth). The *only* way to find The One Great Truth, they believe, is to always follow the Cloud Signs, where the Cloud Signs are interpreted, by the Zazande's sacred book, into principles which guide their judgments about daily affairs. The Cloud Signs (as interpreted by the sacred book) are of course unreliable. But raising this point does not dissuade the Zazande; they believe that the Cloud Signs afford good principles because they will bring one to learn the One Great Truth – something they are convinced is gained only after many years of following the Cloud Signs.

To appreciate why overcoming bootstrapping-generated circularity fails to *de facto* overcome the variety of circularity that's at the heart of incommensurability arguments, we can simply consider (with reference to the Zazande story) how the variety *at issue* in incommensurability arguments might persist even were the variety at issue in bootstrapping arguments to be *ex hypothesi* overcome. In particular, just imagine that (never mind how) an epistemic angel provides a group of Westerners with a kind of 'divine knowledge' that their epistemic principles were *reliable* – such that the group of Westerners would have no need whatsoever to appeal to their own epistemic principles in order to establish this.<sup>32</sup> The group of Westerners' bootstrapping problem, let's suppose, is overcome. They now, let's suppose, non-circularly know that their epistemic principles are reliable – and they also (let's suppose further) are – *unlike* Norman the clairvoyant – able to articulate their grounds: the visit from the epistemic angel.

However, their unique situation (in encountering the Zazande) is one where the Zazande could very well be completely happy to grant that their (the Westerners') principles are reliable. However the Zazande still request *further* reason why the Westerner's system (divinely established as reliable) is worthy of adopting (for they find reliability something of a red herring). Reliability, recall, is not important to the Zazande. The Westerner then is in then back in the original position: in this case, despite knowing (*ex hypothesi*, non-circularly) that her belief forming processes are *reliable*, she must inevitably apply a system within which reliability is regarded as epistemically valuable in order to establish that reliability is epistemically valuable to the Zazande for whom this is in

doubt. Thus, as this example with the Zazande is meant to show, the variety of circularity that's really at issue can't be stomped out *simply* were one to offer a satisfying response to the kind of circularity that affects the BKEP-style reliabilist-cum-boostrapper.

### 5.3.2 Varieties of application circularity

The remarks in the previous section constitute, taken together, a presumptive case for thinking that an externalist route to blocking the kind of application circularity that seems to arise from the original position, is not by any means a simple and direct one. I make no claim about whether an externalist solution could be *ultimately* satisfactory – either toward the aim of stomping out (or undercutting) circularity, or, more generally, for vindicating non-relativist objectivity about our epistemic principles.

Let's now, bearing this in mind, refocus on the question about the nature of the epistemic circularity that *does* arise, when we attempt to justify our own epistemic principles, from within our own system. One immediate observation about applying one's epistemic system in the service of justifying that system is *normative* – and specifically, to do with the connection between such circularity and epistemic justification. As Richard Fumerton (1995, 180) has remarked:

[...] there is no philosophically interesting notion of justification or knowledge that would allow us to *use* a kind of reasoning to justify the legitimacy of using that reasoning.<sup>33</sup> [my italics]

If Fumerton's right about this, then application circularity can be understood at least in part in terms of its essentially violating some justificatory norm. Put another way, we can think of Fumerton's objection to application circularity as highlighting a restriction on the conditions under which *justification* for a premise can transmit to a conclusion, and then to suppose that something about the structure of application circularity violates that restriction.

But application circularity can plausibly be realised in different ways, in practice, and it's not obvious that all forms of application circularity will be problematic to the same extent in their *justificatory structures*. Let's now look at some of the ways application circularity can be realised – when one (put intentionally roughly) *applies one's own epistemic system in the service of justifying it*.

Call *the justifier* what one applies when justifying something – call this latter thing the *justified*. In a ham-fisted case of application circularity,

the justifier and the justified are the very *same* epistemic principle. For example:

- (11) S applies epistemic principle X of epistemic system  $\phi$  in an attempt to justify epistemic principle X of epistemic system  $\phi$ .

But one might as well apply one rule within a system to justify *another* non-identical rule within the very same system, in such a way that the justifier and the justified differ, but are both principles belonging to the same system:

- (12) S applies epistemic principle X of epistemic system  $\phi$  in an attempt to justify epistemic principle Y of epistemic system  $\phi$ .

Just as well, we can envision cases where the justifier/justified is construed as the epistemic system *itself*, where the system is either meant to justify, or be justified by, a principle within the same system. Consider the following two cases:

- (13) S attempts to justify epistemic system  $\phi$  by applying epistemic principle X of epistemic system  $\phi$ .  
 (14) S attempts to justify principle X of epistemic system  $\phi$  by applying epistemic system  $\phi$ .

An instance of (13) will be when someone is attempting to justify, say, 'Western Science' and applies the epistemic principle we called observation (O), a principle that belongs to the epistemic system 'Western Science', in an attempt to do so.

But how might one even *attempt* to accomplish the task in (14)? As epistemic systems are composed of epistemic principles, one could trivially accomplish (14) in a situation where one's epistemic system  $\phi$  had only one rule: principle X. In the limiting case where one is working with a one-rule system, (14) collapses into (11). But, what if  $\phi$  has 100 rules? Even if someone were smart enough to *apply* 100 rules at one time (and *de facto* apply  $\phi$ ), this wouldn't really describe what one is *doing* when one is in a situation like (14) – *viz.*, where one takes the system as a whole *as* the justifier. Typical cases featuring such a move will be ones where, say, an individual attempts to justify epistemic principle (O) by applying the system, not by *using all of the rules of the system all at once*, but by *using the credentials of the system* as a means of establishing the credentials of a rule within the system.

This notion of applying a rule is familiar in a legal context. A judge might apply a rule (consider, the rule ‘One must: drive only with a license’) not by following the rule but by invoking its authority.<sup>34</sup> Accordingly, as MacCallum (1966) has argued, we can think of broadly two ways one might *apply* a rule: by *following the rule* or by *invoking the authority* of the rule. And sometimes, in the latter case, invoking the authority of a rule might well be done, as McCallum puts it, ‘in support of a course of action already determined.’ (For example, I might invoke the authority of the Bible to justify a course of action I’ve already decided to do, purely out of self-interest.)<sup>35</sup>

The distinction between these two very different ways of applying an epistemic rule – by following the rule and by invoking the authority of the rule – mark two interestingly different *species of application circularity*, where application circularity is the ‘genus’ whereby (at least) part of an epistemic system is applied in the service of its own defence. We can accordingly distinguish two varieties of application circularity, one variety to correspond with each species: call these *rule-following application circularity* and *invoking application circularity*.

With respect each of these distinct species of application circularity, we can then ask more fine-grained questions: in the case of rule-following application circularity (for short: *rule circularity*), we can ask: is one following a rule within an epistemic system  $\phi$  in an attempt to justify that very same rule  $\phi$ ? Or, is one following a rule within an epistemic system  $\phi$  in an attempt to justify a *different* rule,  $Y$ , within system  $\phi$ ? Or, is one following a rule within  $\phi$  in order to justify  $\phi$  (as a whole)?

Correspondingly, in the case of *invoking application circularity* (for short: *invoking circularity*), we can ask: is one invoking (the authority of) an epistemic principle  $X$  within an epistemic system  $\phi$  in an attempt to justify that very same epistemic principle  $X$  of epistemic system  $\phi$ ? Or, is one invoking an epistemic principle  $X$  within an epistemic system  $\phi$  in an attempt to justify a different epistemic principle,  $Y$ , within system  $\phi$ ? Or, is one invoking an epistemic principle  $X$  within  $\phi$  in order to justify  $\phi$ ? And finally, in the case of invoking application circularity – though not in the case of rule-circularity – we can ask yet a further question: is one invoking (the authority of) an epistemic *system*  $\phi$  in order to justify that very epistemic system,  $\phi$ , or an epistemic principle within  $\phi$ ?

While this looks admittedly a bit messy, there is I think a reasonably elegant way to assess the epistemic *significance* of both forms of application circularity, rule-following and invoking, under the same general rubric.

Here it is helpful to take a cue from Paul Boghossian (2001, 23–24), where the context of his remark is an assessment of the justificatory structure of arguments that take the rule-circular pattern. He writes:

If rule-circular arguments are in fact capable of transferring warrant from their premises to their conclusions, we should expect this result to flow in some natural way from the conditions that govern warrant transfer *quite generally*.<sup>36</sup>

This idea has some promise. Why not think that, if rule-circular arguments have a defective justificatory structure, that we should be able to *explain this* with reference to some more *general* explanation for why *any* piece of reasoning's justificatory structure is defective?

While Boghossian himself uses Alston's (1986) general account on this score as a reference point for defending the justificatory structure of some kinds of rule-circularity, I think that the justificatory structure of rule-circular *and* invoking-circular species of epistemic circularity can be more helpfully subsumed (or so I'll shortly argue) within the purview of what Jim Pryor calls Type 4 and Type 5 argument patterns.

Arguments that feature what Pryor terms Type 4 and Type 5 dependence have been discussed most notably in the literature on perceptual warrant. And the situating reference point for such discussions is a topic approached for different reasons in Chapter 1 in the section on Moore's Proof.

### 5.3.3 Application circularity and cogency

Consider the following piece of reasoning, due to G.E. Moore (1939):

15. Here are two hands.
16. If hands exist, then there is an external world.
17. So there is an external world.

As Pryor (2004, 349) observes, 'Something about this argument sounds funny.' But what is it? More specifically: what's fishy about purporting to acquire justification for this conclusion by reasoning through these premises – *even though* the premises are ones we take ourselves to know, *and* they entail the conclusion? Suitably *general* answers to this question should shed some light on the question of why application-circular arguments (of *either* species: rule-circular as well as invocation-circular arguments) might be regarded as failing to transmit warrant from premise to conclusion. Here's a very general remark to make: it seems like Moore's

first premise in some way *depends* on the conclusion. Though from here things get thorny.

Not all ways in which a premise depends on a conclusion should be regarded as objectionable. As Pryor (2004, 359) has pointed out, whatever else we want to say is wrong with the reasoning that features in Moore's proof, it *can't* be that, in order to have justification to believe Moore's premise (here are two hands), it's *necessary* that you have justification to believe the conclusion (there's an external world). As Pryor (2004, 359) writes, this kind of dependence (which he terms 'Type 3')

[...] seems to include some arguments that are perfectly respectable. It includes arguments where the connection between premise and conclusion is so obvious that understanding the premise well enough to be justified in believing it *requires* you to take any justification for the premise to also justify you in believing the conclusion.

As Pryor sees it, here's where the real action lies: between what he calls Type 4 and Type 5 Dependence.

**Type 4 Dependence:** the conclusion is such that evidence against it would (to at least some degree) undermine the kind of justification you purport to have for the premises.

**Type 5 Dependence:** having justification to believe the conclusion is among the conditions that make you have the justification you purport to have for the premise.

Following Pryor, we can demarcate the key dividing line in contemporary debates about perceptual warrant *with reference to* whether Moore's Proof exhibits Type 4 or Type 5 dependence.

Here's what the *conservativists* (e.g. Wright (2007); Davies (2003)) say – at least, the bit they say which is relevant for now. The conservativist line, as Pryor (2004) sees things, is that Moore's Proof exhibits *Type 5 Dependence* – *viz.*, that Moore's Proof fails to transmit warrant from premise to conclusion because *any Type 5 argument fails to transmit warrant from premises to conclusion*, where this claim about Type 5 dependence, as such, is being put forward as a *general* thesis about warrant transmission.<sup>37</sup>

*Dogmatists* about perceptual warrant diagnose Moore's Proof differently. In doing so, they reveal another interesting general thesis that will be useful in what follows. Pryor, himself a dogmatist, agrees that

'Type 5 dependence ruins an argument' (*ibid.* p. 360) – which by this he means that Type 5 arguments have a defective justificatory structure: warrant for the premises does not transmit to the conclusion. However, as Pryor sees it, it is possible for an argument to exhibit *Type 4* dependence while failing to exhibit Type 5 dependence.<sup>38</sup> Secondly, arguments that exhibit *merely* Type 4 dependence are not 'epistemologically objectionable' – that is, warrant for believing the premise *can* (provided the argument is merely Type 4, and not Type 5) transmit from premise to conclusion; and thirdly, Moore's argument, according to Pryor, exhibits Type 4 and *not* Type 5 dependence. In general terms, Pryor's view is that, while *Type 5 arguments fail to transmit warrant from premise to conclusion*, Type 4 Arguments don't, but – and this is an important concession – *Type 4 Arguments are dialectically ineffective* against one who antecedently doubts the conclusion.<sup>39</sup> That is, Pryor concedes that all Type 4 arguments are such that – while there's nothing wrong with their justificatory structure as such – they are not effectual in bringing one who *already* doubts the conclusion to rational conviction of the conclusion on the basis of the premises.

A quick note to the reader: the matter of who's *right* between Pryor and Wright won't matter for our present purposes. What matters is that we have *extracted* from their debate two entirely general theses about warrant transmission. They are also, interestingly, both general theses which Wright and Pryor, *despite their differences*, can accept: these are that

- *Type 5 arguments* fail to transmit warrant, and;
- *Type 4 arguments* are dialectically ineffective.

### 5.3.4 Application circularity revisited

The aim of this section will be outline how the two species of the genus application circularity (e.g. rule-following and invoking application circularity) can interface with – specifically, by being subsumed (modulo certain tweaks) *within* – the Type 4 and Type 5 structure. Then, in the final section, I'll suggest what the upshot of this is (with respect to the plausibility of the move from application circularity to epistemic relativism) by drawing out some further parallels to the perceptual warrant debate, and in doing so, outline an answer to Q3.

Let's return now to the distinction made between the two primary varieties of application circularity, *viz.* *rule-following* application circularity and *invoking* application circularity, focusing on the *invoking* strand first.

Suppose (from the original position), I attempt to justify the principle Boghossian calls *Observation* (O). For ease of reference:

*Observation* (O): For any observational proposition  $p$ , if it visually seems to S that  $p$  and circumstantial conditions D obtain, then S is *prima facie* justified in believing  $p$ .

In the case of the genus *invoking application circularity*, we can imagine *three* versions of how this might go. On the first version (corresponding with (11)) I might try to justify an epistemic principle by invoking the *authority* of that very principle – *viz.*, by *appealing to* (O) in attempting to establish that (O) is justified. If in a belligerent mood, I might attempt to do this by reasoning like this: O, therefore, O. This trivially exhibits both Type 4 dependence *as well as* Type 5 dependence.

But I might also appeal to O in an attempt to establish that O is justified in the following way, which is a slightly less blatant way to appeal to O in the service of establishing O. Let  $\alpha$  be an arbitrary proposition.

18.  $\alpha$  (on the basis of O)
19. If  $\alpha$ , then O
20. O

*Question:* how should this strand of invoking application circularity interface with the Type 4/Type 5 distinction? I submit that such a line of reasoning (appealing to the authority of O in order to establish O) will exhibit *at least* Type 4 dependence: evidence against O would (to at least some degree) undermine the kind of justification you purport to have for the premise (O). As such, this argument will be dialectically ineffective. Perhaps, the move from 18 to 20 also exhibits Type 5 dependence (though as I'll suggest later, this won't obviously be agreed upon by dogmatists and conservativists). In sum: definitely Type 4; maybe Type 5.

Let's consider a slightly different case of invoking application circularity. I might try to justify (O) in a way that corresponds with (12): that is, I might attempt to justify (O) by invoking the authority of *another* non-identical epistemic principle belonging to the same system – say, Inference to the Best Explanation (IBE). For instance:

21.  $\alpha$  (on the basis of IBE)
22. If  $\alpha$ , then O
23. O

It's not obvious that the move from 21 to 23 would exhibit Type 5 dependence (after all, it's not obvious that you are justified in believing some proposition on the basis of IBE in part in virtue of justification you already have for O). *However*, it's hard to see how such a move would *not* exhibit Type 4 dependence.<sup>40</sup> And, finally, the third version of invoking application circularity will be one where I attempt to justify (O) not by invoking the authority of (O) itself, or even by invoking some other principle within the wider epistemic system I embrace – but rather, by invoking the authority of that *system*. This version is a bit more complicated. Assume 'Western Science' is an epistemic system that includes (O), (IBE) and other basic epistemic principles and their corresponding rules, which are typically taken for granted. Now, suppose I invoke the authority of Western Science (*qua* system) in the service of establishing the epistemic credentials of (O). Here's a natural way this might go:

24. Any epistemic principle indispensable to Western science is *prima facie* epistemically justified.
25. Epistemic principle (O) is indispensable to Western science.
26. Therefore, (O) is *prima facie* epistemically justified.

*Diagnosis*: this strand of reasoning from 24–26 clearly exhibits Type 4 dependence. Grounds to doubt (O) and thus (26) are at the same time grounds to doubt the support you take yourself to have for (24). So 24–26 are dialectically ineffective – *viz.*, you couldn't bring one who already doubted (O) to rationally accept (O) on the basis 24 and 25. And it *might* as well be (though this is less clear) that attempting to justify a given epistemic principle by invoking the authority of the system within which that principle is a component exhibits Type 5 dependence.<sup>41</sup>

Summing up this evaluation of *invoking* application circularity: the invoking species of application circularity is *always* dialectically ineffective, and at least sometimes fails to transfer warrant. That at any rate is the result of mapping out the varieties of invoking application circularity in terms of the patterns which individuate Type 4 and Type 5 cases.

Let's turn now to the *rule-following* species of application circularity. For simplicity and easy of reference, let's follow Boghossian's (2001) example that features in his discussion of rule-circularity. Let's suppose I am attempting to *justify* modus ponens.<sup>42</sup> How might I do this? Suppose I try to produce an argument, the conclusion of which is modus ponens.

Let  $\alpha$  be some consideration I appeal to in an attempt to argue inferentially for MP:

- 27.  $\alpha$
- 28.  $\alpha \rightarrow \text{MP}$
- 29. MP<sup>43</sup>

Boghossian, in considering what would be needed for such an argument to have a satisfactory justificatory structure, quickly draws attention to something that he regards *must not* be necessary (lest we be sceptics), which is a condition Alston (1986) had thought was necessary for a piece of reasoning to confer justification to the conclusion. Alston's condition was one must know, or justifiably believe, that the premises and conclusion are logically related in such a way that if the premises are true, that is a good reason for supposing the conclusion is likely to be true.<sup>44</sup>

In short, Boghossian rightly worries that insisting on *this* condition is a one-way ticket to a Lewis Carroll-style regress.<sup>45</sup> For if such a condition *were* in place, then warrant could transmit from 27–29 *only if* one is already justified in believing that the premises necessitate the conclusion. And such justification would itself (on a model where this condition is embraced) depend on further justification, *ad infinitum*. One lesson that the Lewis Carroll (1895) regress has suggested – one which has been famously exploited by Gilbert Ryle (1945) in the service of arguing that knowledge how is grounded in abilities rather than knowledge of facts – is that (as Boghossian puts it):

at some point it must be possible to use a rule in reasoning in order to arrive at a *justified* conclusion, without this use needing to be supported by some knowledge about the rule that one is relying on. It must be possible to simply move between thoughts in a way that generates justified beliefs, without this movement being grounded in the thinker's justified belief about the rule used in the reasoning. (*Ibid.*, 27)

The situation seems to be this: *unless* one wants to defend a very strong intellectualist model of basic mastery of logical rules, and then explain how such a model sidesteps the Carroll regress, then it does seem that to the extent that justification for the basic rules underwriting our epistemic principles can be acquired at all, this needn't be *via* (Alston-style) justified beliefs about the legitimacy of moving 'between thoughts' but can also come by simply moving through thoughts in a way that generates justified beliefs.

Now, if the view that Type 4 arguments are dialectically ineffective and Type 5 arguments fail to transmit warrant really are thoroughly general theses about justificatory structure, then we really ought to be able to think of what's going on in rule-circular cases (as when we take a step in accordance with modus ponens in order to establish modus ponens) *in terms of* exhibiting (some variations on) Type 4 and/or Type 5 dependence. Put another way: we ought to be able to *formulate* variations on Type 4 and Type 5 dependence in a way that naturally subsumes varieties of rule-circular argument. And in doing so, what will be relevant to whether a piece of reasoning exhibits Type 4 or Type 5 dependence is not whether and to what extent justification for the premise depends in some way on the conclusion. Rather, what will be relevant is whether the *legitimacy of employing the rule one does when moving from premises to conclusion depends in some way on the conclusion itself*.

There's a pretty straightforward way to model Type 4 dependence along these lines:

**Type 4-Rule Dependence:** the conclusion is such that evidence against it would (to at least some degree) be evidence against the legitimacy of employing one (or more) rules one reasons in accordance with in moving from premises to conclusion.

And just as Type 4 dependence arguments are dialectically ineffective, so will be Type 4 Rule circular arguments. (For example: one who already doubts modus ponens won't be rationally persuaded to accept it on the basis of an argument that reasons in accordance with it.) Type 5-Rule Dependence is (considerably) trickier to model. Here's an attempt:

**Type 5-Rule Dependence:** The legitimacy of reasoning in accordance with the (content of the) conclusion just is (or, is partly constitutive of) the legitimacy one purports to have for moving from premises to conclusion.

Suppose now that one attempts to make the move from 27–29 (where  $\alpha$  is arbitrary, as it features in 27), and where clearly, one is reasoning in accordance with MP, in the sense that, one takes *at least one step* in accordance with MP. What's amiss with such a piece of reasoning? The first, and obvious, observation is that the move from 27–29 exhibits *Type 4-Rule dependence*; evidence against the conclusion would (to at least some degree) be evidence against the legitimacy of employing one (or more) rules one reasons in accordance with in moving from

premises to conclusion. Thus, the move from 27–29 will be dialectically ineffective (as *per* any Type 4 argument). And indeed, that's exactly what Boghossian thinks – he remarks that:

if we were confronted by a skeptic who *doubted* the validity of [MP] in any of its applications we could not use [a Type 4-Rule Argument] to rationally persuade him. (*Ibid.*, 35)

For one who takes the ability to rationally persuade the sceptic as criterial for our own justification, that fact the reasoning in 27–29 exhibits Type 4 Rule-dependence is troubling. But Boghossian thinks that since scepticism is simply unacceptable, we should just *deny* that the ability to rationally persuade the sceptic should be regarded as criterial of our own justification. In short, we cannot have what Boghossian calls *suasive* reasons for MP, though we can have *non-suasive* reasons (where non-suasive reasons are powerless to rationally quell sceptical doubts).

While Boghossian's right that taking at least one step in accordance with modus ponens is enough to leave an argument *for* modus ponens dialectically ineffective – *viz.*, on our taxonomy, a Type 4-Rule Dependent argument – Boghossian might potentially be too optimistic in stopping there. Perhaps an argument for modus ponens that reasons in accordance with modus ponens also exhibits *Type 5-Rule* dependence, in which case (by analogy to Pryor's Type 5 dependence) warrant *wouldn't* transfer to the conclusion<sup>46</sup>. If this were so, we would lack the ability to acquire *suasive* *or*, and contrary to what Boghossian thinks, *non-suasive* reasons by applying that rule in an argument attempting to prove it. However, just as one might apply one rule within a system to justify another rule, by *invoking* that rule, so also one might attempt to apply one rule within a system to justify *a different rule* within that system, by reasoning in accordance with it. One might reason in accordance with modus ponens in order to prove IBE, or infer MP as an inference to the best explanation. In such a case (by parity of reasoning, from the case of invoking application) we'd have Type 4 rather than Type 5 dependence.

## 5.4 Some parallels

The key result from the previous section is a certain observation:

- (i) The kinds of epistemic circularity that in fact arise when we attempt to justify *either* our epistemic system as a whole or any of our individual epistemic principles by applying (*either*: by invoking

the authority of, or by reasoning in accordance with) our own epistemic principles will exhibit the structural features of *either* (i) Standard Type 4 Dependence or Type 4-Rule Dependence; or (ii) Standard Type 5 Dependence or Type 5-Rule Dependence.

And here is a second observation:

- (ii) Though Wright and Pryor disagree fundamentally about perceptual warrant, they agree that: *if* a given piece of reasoning is a Type 4 case, then it is dialectically ineffective (incapable of rationally persuading one who antecedently doubts the conclusion); and *if* a given piece of reasoning exhibits Type 5 dependence, it fails to transmit warrant from premise to conclusion.

And, now, consider a third claim, the ramifications of which I want to now explore, in the service of addressing Q3 – regarding the move from epistemic circularity to epistemic relativism:

- (iii) Wright and Pryor propose differing *anti-sceptical* strategies aimed at vindicating perceptual warrant in the face of the threat of epistemic circularity. The three central options they regarded as on the table are (i) dogmatism, (ii) conservatism; and the unacceptable option, (iii) scepticism; neither chooses scepticism, and relativism isn't even on the radar.

In what follows I want to briefly sketch the dogmatist and conservativist *vindicatory* (i.e. anti-sceptical) strategies. I want to then suggest that the most salient options on the table – in the face of the Type 4 and/or Type 5 dependence that arises off the back of application circularity – should be, by parity of reasoning, *analogues to the dogmatism, conservatism and scepticism* positions *vis-à-vis* perceptual warrant. My move for blocking the circularity-to-relativism leg of the incommensurability–circularity–relativism sequence is thus an argument from parity. Put another way: incommensurability style arguments for epistemic relativism which rely on epistemic circularity in generating epistemic relativism motivate epistemic relativism only in so far as locating the same varieties of epistemic circularity in the perceptual arena should motivate epistemic relativism *over* one of the three salient options of dogmatism, conservatism or scepticism. Relativism isn't motivated by epistemic circularity in the perceptual case and by parity of reasoning we have no reason to think it should in the case of epistemic principles either<sup>47</sup>.

Let's return to Moore's proof as a frame of reference – this time, to remark upon how Wright and Pryor distinguish their *anti-sceptical vindicatory strategies*.

30. Here are two hands.

31. If hands exist, then there is an external world.

32. So there is an external world.

*Dogmatists* about perceptual warrant maintain that you are justified in believing (30) on the basis of the perceptual evidence as of seeing hands *provided* you have no reason to doubt (32). And if that's the case, there's no 'transmission failure' in the argument. The only problem is that the argument is *dialectically ineffective*. Dogmatism is thus *vindicatory* in the face of Type 4 circularity: dogmatists allow one to attain warrant for the conclusions of Type 4 arguments.

By contrast *conservativists* about perceptual warrant tell a different kind of vindicatory story. On Wright's view, Pryor is wrong that you are justified in believing (30) on the basis of perceptual evidence *provided* you lack any antecedent doubts. On the conservativist line, one is justified in believing (30) on the basis of the perceptual evidence only if one has positive independent reasons or warrant for believing (32). This is sounding pretty sceptical but Wright *does* think we have such warrant, in the form of what he calls *entitlements*. It is in virtue of our entitlement to (32) that we are justified in believing (30) on the basis of the evidence. However, it follows from this entitlement line that one's warrant for (32) is also part of what warrants one in accepting (30) on the basis of the perceptual evidence of the appearance of a hand. And so an implication of the conservativist's vindicatory line is that the piece of reasoning above exhibits Type-5 dependence. And accordingly, Wright denies that the warrant for the premise (warrant that on the conservativist line is warrant that consists in the entitlement to (17)) transmits from premises to conclusion. Both Wright and Pryor think their vindicatory lines are needed to avoid scepticism.

Let's imagine now *analogous* vindicatory lines, against the threat of epistemic circularity, in the case of *epistemic principles*. Call *dogmatism-EP* the view (analogous to dogmatism about perceptual warrant) that (put roughly for now) you are justified in believing the *deliverances* of our epistemic principles (e.g. what is entailed, in particular cases, by the *application* of our epistemic principles) provided we have no reason to doubt that our epistemic principles are correct. We can define deliverances in a way that maps on to the two ways we can *apply* epistemic

principles: by invoking their authority and by reasoning in accord with them. So, dogmatism-EP says that we are justified in believing the deliverances of our epistemic principles (understood in either the invoking or rule-following sense) provided we have no reason to doubt these principles. And moreover, to continue the analogy, dogmatism-EP says that warrant *can* transmit across a piece of reasoning that concludes that some epistemic principle is true and which relies on – as a premise in the piece of reasoning – the deliverances of our epistemic principles *provided* we don't already doubt the principle that features in the argument's conclusion<sup>48</sup>.

What we can call *Conservativism-EP*, by contrast, will insist that that we are justified in believing the deliverances (e.g. what is entailed, in particular cases, by the *application* of our epistemic principles) of our epistemic principles provided we *already* have positive reasons or warrant for accepting them. (This is the structural analogue at any rate.) And, since it is in virtue of our entitlements to epistemic principles that we are justified in believing their deliverances, warrant doesn't transmit from premises appealing to the deliverances of our epistemic principles to the conclusion that our epistemic principles are true.

These are, in rough form, the two analogue positions, with the third of course being *scepticism* – the view to which (in the perceptual arena) Pryor and Wright are offering competing antidotes. The note I want to leave things on in this chapter is this: incommensurability arguments put us in a unique position where it is not entirely irrelevant to attempt to justify our own epistemic principles and/or our own epistemic system. At this point, it looks as though it is inevitable that we apply our own epistemic principles in the service of this task. The incommensurability–circularity–relativism sequence regards the kind of circularity at issue as circularity that leaves all epistemic principles on an even footing – more specifically, those who attempt to generate epistemic relativism off the back of, for example, Kuhn-style, Wittgenstenian-style and Hacking-style incommensurability rely on the circularity that materialises once one *applies* one's epistemic principles in the service of defending them to *motivate* epistemic relativism.

But as I've argued, closer attention to the kind of application circularity that follows in the wake of attempting to justify your own principles doesn't motivate epistemic relativism *any more than*, or so I've suggested, the structurally same kind of circularity motivates epistemic relativism in the perceptual arena – an arena where the salient options are dogmatism, conservativism and scepticism<sup>49</sup>. To make this point, I showed how application circularity comes in two forms: invoking

and rule-following. I showed how invoking application circularity will exhibit a structure of *either* what Pryor calls Type 4 or Type 5 dependence; further, I showed that rule circularity will involve *either* Type 4 or Type 5 Dependence.

Pryor and Wright, in the face of the threat of epistemic circularity that takes a Type 4 or Type 5 form, each propose a *vindictory* strategy, dogmatism and conservatism – each of which is regarded (respectively) as necessary to avoid scepticism. I've argued that, since application circularity is Type-4 or Type-5, we are in the relevantly same position: we can envision vindictory strategies that line up with dogmatism and conservatism, respectively, with respect to the deliverances of the application of our epistemic principles. Each vindictory strategy will then (like Wright's and Pryor's) regard itself as necessary for avoiding scepticism. And having drawn the parallel as such, I submit that having suitably uncovered the nature of application circularity, we have reason to think that relativism is recommended *only if* perceptual relativism is motivated by the structurally same variety of circularity that arises off the back of Type 4 and Type 5 arguments whereof one appeals to the deliverances of perception rather than the deliverances of one's epistemic principles. Perceptual relativism is not really a live option in this debate, and so it is incumbent on the incommensurability-style epistemic relativist to tell us why the salient options should be anything other than what we should expect: dogmatism, conservatism and scepticism.

# 6

## Replacement Relativism: Boghossian, Kusch and Wright

*Abstract.* Even if epistemic relativism could be philosophically motivated, there remains the issue of how to formulate it. This chapter engages with the plausibility of one popular semantic strategy for making sense of epistemic relativism: the *replacement relativism* model (e.g. Boghossian 2006a; 2006b; Kusch 2010). I begin by outlining how Boghossian thinks of epistemic relativism along the lines of the replacement model, and as such regards the view as incoherent. I then consider in some detail a recent and creative attempt by Kusch (2010) to defend a version of epistemic relativism against Boghossian's criticisms within the replacement model that Boghossian thinks is unworkable, in particular, by defending an 'incompleteness-theoretical' version of the view which differs in important ways from the strand that Boghossian attacked under this description. I raise several worries for Kusch's strategy but argue that, even if these worries can be overcome, a difficult issue remains: at the end of the day, Boghossian and Kusch will, by Wright's (2008) lights at least, both be tarred with the brush of not formulating the epistemic relativist's core position in an appropriately charitable way. And this charge, I shall suggest, is not one that can be easily dismissed. On the interpretation of this standoff that I shall propose, we'll come to see that the core idea driving Burnyeat's (1976) *principle of translation* – a principle Burnyeat (1976) regarded as essential to making sense of Protagorean subjective global relativism – will re-emerge as an insight about relative truth that is, I think, central to what is fundamentally at issue between Wright on the one hand, and Boghossian and Kusch the other. After briefly exploring some further avenues out of the deadlock, I conclude by connecting the state of the debate about replacement relativism with a very different semantic

model for representing epistemic relativism – John MacFarlane’s (2005, 2011, 2014) *assessment-sensitive* model – which will be the central topic of the next two chapters.

## 6.1 Introduction

The past three chapters have pointed to a conclusion: familiar rationales for epistemic relativism come up short. At least, the Pyrrhonian, dialogic, and incommensurability argument-types were shown to fail to *distinctively* motivate epistemic relativism over other available alternatives, particularly, scepticism. But let’s suppose for a moment that the bulk of the objections raised over the past three chapters could be overcome, or perhaps, that there is some further persuasive argument which indicates that epistemic relativism is true.

Even on this supposition, there remains another issue, which is just how to *formulate* the position in a coherent manner. And this brings us back to a thorny issue that was raised, but then bracketed, in Chapter 2. Recall that when Burnyeat (1976, 192–93) was thinking about how to make *sense* of subjective global relativism (of a Protagorean sort) he suggested that at the end of the day, we simply must attribute to the Protagorean relativist what Burnyeat called the *principle of translation*, according to which, a proposition of the form ‘x is F’ is true (relatively) for person (a), if and only if, ‘x is F for a’ is true (absolutely).<sup>1</sup> I won’t rehearse here my assessment of Burnyeat’s move, in the context of formulating global relativism. Rather, I want to remind the reader that I noted that the matter of whether making *sense* of global relativism *requires* that we attribute to her a principle of this sort foreshadows a more contemporary dispute between Boghossian and Wright about how to best formulate *epistemic* relativism.

The battleground for this more contemporary debate surrounds the viability of *replacement relativism* as a semantic strategy for formulating philosophical forms of relativism, such as epistemic relativism. As Martin Kusch (2010, 165) rightly notes, replacement relativism is the ‘main’ semantic strategy for representing what the epistemic relativist is putting forward about first-order epistemic judgments. The core idea at play featuring in the ‘replacement relativist’s semantic strategy is that, as Boghossian (2006b, 21) puts it, we should understand the relativist as aspiring to ‘replace’ what’s ‘doomed to falsehood’ (e.g. unrelativised first-order judgments of the form ‘X is F’) with the nearest absolute

truths in the vicinity which the relativist can accept, which are explicitly relational judgments (e.g. X is F relative to parameter P).<sup>2</sup>

In this chapter, I'll begin by showing how it is Boghossian first formulates the epistemic relativist's position within the replacement model and then argues that epistemic relativism is ultimately *incoherent*. Next, I consider in some detail a recent and creative attempt by Martin Kusch (2010) to defend a version of epistemic relativism within the replacement model.<sup>3</sup> I raise a few worries for Kusch's strategy but argue that, even if these worries can be overcome, an issue remains which is that – at the end of the day – Boghossian and Kusch will, by Wright's lights, both be tarred with the brush of not formulating the thesis in a satisfactory way. As I shall set things up, the core idea driving Burnyeat's (1976) principle of translation will re-emerge as central to what is at issue between Wright on the one hand, and Boghossian and Kusch on the other. I conclude by connecting the state of the debate about replacement relativism with a different semantic model for representing epistemic relativism – John MacFarlane's (2005, 2011, 2014) *assessment sensitive* model – which will be the central topic of the next two chapters.

## 6.2 Boghossian on replacement relativism

In Chapter 6 of *Fear of Knowledge*, Boghossian (2006a) goes for the jugular: the conclusion of this chapter is that there is 'no way of construing the notion of an epistemic system so as to render stable a relativistic conception of epistemic justification' (Boghossian (2006a, 92–93). Let's consider how he reaches this conclusion.

Boghossian begins by considering a particular unrelativised epistemic judgment and then details what the epistemic relativist is, by her own lights, supposed to be telling us about it. The judgment Boghossian begins with is this particular judgment:

1. Copernicanism is justified by Galileo's observations.

As Boghossian sees it, the relativist views (1), and all judgments like it (e.g. particular unrelativised epistemic judgments) as 'doomed to falsehood' and the reason is that by the relativist's lights there are no *absolute* facts about justification.

So how is the relativist meant to retain epistemic discourse, then, if (1)-style statements are all false? Boghossian (2006a, 84) remarks:

[...] the relativist urges, we must reform our talk so that we no longer speak simply about what is justified by the evidence, but only about what is justified by the evidence according to the particular epistemic system that we happen to accept, noting, all the while, that there are no facts by virtue of which our particular system is more correct than any of the others.

What Boghossian regards as the core epistemic relativist thesis – betrayed in the rationale quoted above – is unpacked as follows, as a conjunction of *epistemic non-absolutism*, *epistemic relationism* and *epistemic pluralism*.

***Epistemic Relativism (Boghossian’s Formulation<sup>4</sup>)***

- A. There are no absolute facts about what belief a particular item of information justifies. (*Epistemic non-absolutism*)
- B. If a person, S’s, epistemic judgments are to have any prospect of being true, we must not construe his utterances of the form  
 “E justifies belief B”  
 as expressing the claim  
*E justifies belief B*  
 but rather as expressing the claim:  
*According to the epistemic system C, that I, S, accept, information E justifies belief B. (Epistemic relationism)*
- C. There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (*Epistemic pluralism*)

Boghossian’s epistemic relativist then is going to insist that we assert not (1), but rather:

- 2. Copernicanism is justified by Galileo’s observations relative to a system, Science, that I, the speaker, accept.

On this way of articulating the core epistemic relativist’s insight, the relativist is effectively asking us to *replace* (1) with (2) – viz., to replace something that must be regarded as absolutely false with the closest candidate in the neighbourhood which, by the relativist’s own lights, can be regarded as absolutely true.

In order to appreciate the implications Boghossian takes this replacement (of 1-type statements with 2-type statements) commitment to have, it will be helpful to sharpen a few further definitions. According to Boghossian, an *epistemic system* is a set of *general* normative propositions – viz., *epistemic principles*. An example, *Observation*, familiar from the previous chapter, says:

(Observation) For any observational proposition *p*, if it visually seems to *S* that *p* and circumstantial conditions *D* obtain, then *S* is *prima facie* justified in believing *p*.<sup>5</sup>

And epistemic principles, such as (O) – the constituents of epistemic systems – are, as Boghossian writes:

just more *general* versions of particular epistemic judgments. They, too, are propositions stating the conditions under which a belief would be absolutely justified, the only difference being that they do so in a very general way and without adverting to particular beliefs held by particular agents at particular times and under particular evidential conditions. (*Ibid.*, 85)

Against this background, Boghossian highlights the following key parallel: if epistemic relativism demands that particular epistemic judgments are categorically false, then given that epistemic principles are just generalised versions of particular epistemic judgments, it looks like epistemic principles, the very constituents of epistemic systems, must *also* be regarded by the epistemic relativist as false.<sup>6</sup> But if *this* is right, Boghossian thinks, a kind of incoherence problem materialises. In short, the problem is that to

make sense of the idea that Galileo thinks he has a *relative* reason for believing Copernicanism while Bellarmine thinks he has a *relative* reason for rejecting it [...] it is crucial to the relativist's view that thinkers *accept* one or another of these systems, that they *endorse* one or another of them and then talk about what they do or do not permit. (*Ibid.*, p. 86)

Obviously, the knee-jerk relativist defence against the critique that acceptance of epistemic systems cannot be explained provided epistemic principles are regarded as absolutely false will be to say that such acceptance can be rendered compatible with an acceptance of the

explicitly *relational* formulation of epistemic principles which relativists *can* accept.

But since the principles *themselves* are still regarded as uniformly false, Boghossian isn't convinced one can coherently count as accepting or endorsing them. He remarks that it is

hard to explain why anyone should *care* about what follows from a set of propositions that are acknowledged to be uniformly false. What sort of normative authority over us could an epistemic system exert, once we have become convinced that it is made up of propositions that are uniformly false?<sup>7</sup>

### 6.3 Fictionalism and non-fictionalism

After raising this dilemma, Boghossian immediately backtracks a bit. It's not *really* the case that the epistemic relativist must regard all Type-1 propositions (e.g. propositions that take the form of 'Copernicanism is justified by Galileo's observations') as *false*, per se. It's just that: it is off limits for the epistemic relativist to regard such statements as absolutely *true*. And since it is specifically the *falsity* of Type-1 statements that (by entailment) generates what Boghossian took to be a problematic commitment to the falsity of the *general* false propositions that constitute the epistemic principles framing a system, there's a potential passage out of the jam: by regarding such statements as simply *not* absolutely true.

There are broadly two ways one might go here, depending on whether the components of epistemic frameworks (e.g. general epistemic principles) are taken to be general propositions or imperatives. Boghossian labels the first route 'fictionalist' views. For ease of exposition, let's follow Kusch (2010, 166) and call the other strand 'non-fictionalism' (where on non-fictionalism the propositions at issue are not even truth-evaluable). Boghossian does not think non-fictionalism works, though the reasons won't concern us here.

Rather, I want to bear down on how it is that Martin Kusch (2010, 166) has attempted to defend a strand of *fictionalism* against Boghossian's criticisms. In order to appreciate Kusch's defence, though, we need to distinguish the strand of fictionalism Kusch wants to defend from the strand he does not. The relevant distinction here can be drawn on the basis of whether the original non-relativised 1-style statements are untrue by being *false* or untrue by

being *incomplete*. For ease of reference, let's follow Kusch and call *error-theoretical fictionalism* the view that 1-style unrelativised epistemic judgments are untrue because *false* and *incompleteness-theoretical fictionalism* the view that 1-style unrelativised epistemic judgments are untrue because incomplete.<sup>8</sup>

Just as we saw that Boghossian does not think regarding 1-style statements as absolutely false can underwrite a coherent epistemic relativism, he thinks the same is the case for 'incompleteness-theoretical' fictionalism. Let's now look at incompleteness-theoretical fictionalism more closely.

#### 6.4 Incompleteness-theoretical fictionalism

According to incompleteness-theoretical fictionalism, judgments of the (1)-form – *viz.*, Copernicanism is justified by Galileo's observations – are untrue for the same reason that:

Tom is taller than ...

is untrue – because

it cannot be *evaluated* for truth. It is untrue because incomplete. This suggests an alternate way of formulating epistemic relativism [...] What the relativist has discovered is that [statements like (1)] need to be completed *by reference to an epistemic system* before they can sensibly be appraised for truth.<sup>9</sup>

Boghossian's envisioned incompleteness-theoretical fictionalist replacement epistemic relativist (hereafter, ITFRR) therefore regards 'Copernicanism is justified by Galileo's observations' as an incomplete proposition (as Tom is taller than ... is incomplete),<sup>10</sup> which gets completed like this:

[In relation to epistemic system C], Copernicanism is justified by Galileo's observations.

To reiterate: the main *prima facie* appeal Boghossian envisions for such a position is that while *false* 1-style statements will, by entailment, lead to false epistemic principles, incomplete propositions look like they might circumvent this result.

## 6.5 Contra ITFRR: four arguments

Boghossian thinks that whatever promise ITFRR seems to have (as an alternative to error-theoretical fictionalism) fails to pan out. Boghossian's 'anti-incompleteness' arguments are to be found in both *Fear of Knowledge* and also in a paper 'What is Relativism?',<sup>11</sup> written the same year. Helpfully, Kusch (2010, 167) has cobbled together what he takes to be the four principal problems which 'Boghossian counts as decisive' against ITFRR – problems which Kusch thinks (at least, once ITFRR shows its best face) are not decisive: the *normativity problem*, the *endorsement problem*, the *infinite regress problem* and the *entailment problem*.

The *normativity problem* is, in short, that the unrelativised 1-type statements (e.g. Copernicanism is justified by Galileo's observations) are *normative*. But what they are being replaced with – viz., 'In relation to epistemic system C, Copernicanism is justified by Galileo's observations' – is just a *description* of what is epistemically justified *according to* an epistemic system. In short, replacing the incomplete proposition with the completed proposition *loses the normativity* in the process of the replacement.

The second problem concerns *endorsement*. Just as Boghossian thinks there is an endorsement problem for error-theoretic fictionalists (because: how can we endorse what are by our own admission *false* general principles?), he poses the same question to the proponent of ITFRR: how can we endorse general principles that are by our own admission *incomplete*? And if particular (1-type) epistemic judgments are all incomplete, then since epistemic principles are just general versions of particular epistemic judgments, epistemic principles will be incomplete as well. And it's just not clear how something incomplete can be *endorsed*.

Here's an idea: fill them in, and *then* endorse them! But this idea hints at the *infinite regress problem*: the very specification of an epistemic system that is needed to *complete* an incomplete proposition (whether it be an incomplete particular epistemic judgment, or an incomplete general epistemic principle) involves identifying general epistemic principles. But these are incomplete as well and must be completed, *ad infinitum*.

The fourth key problem for ITFRR is the entailment problem: although the fact that there are entailment relations between particular epistemic judgments and general epistemic principles contributed to some of the problems for error-theoretical fictionalism, at least error-theoretical fictionalism can explain how general principles can entail particular judgments. But it's not clear how such entailment relations can hold insofar as the general principles are supposed to be incomplete.<sup>12</sup>

## 6.6 Kusch's defence of incompleteness – theoretical replacement relativism

It would seem as though incompleteness-theoretical replacement relativism is dead in the water. However, with some tweaks, and some repositioning, this strategy-type might be able to do considerably better than did the *variety* of it which Boghossian attacked. And indeed the variety Boghossian attacked is but one way to think about what incompleteness-theoretical replacement relativism might be.

### 6.6.1 Kusch's revised ITFRR

Kusch's argument has several working parts which, when put together, support an original rationale for ITFRR, one which he thinks can solve all four problems outlined in the previous section.

One such working part involves reconceiving the logical structure of the 'replacing proposition'. Kusch's thinking is that – given that relativism is unobjectionable in the physical case (e.g., Galileo's discovery that motion is relative) – we're best off modelling semantically, to the extent that it's feasible to do so, the relativist's discovery in the *philosophical* case off the relativist's discovery in the physical case.<sup>13</sup>

In this physical case, we can think of Galileo as having shown that statements of the form 'X moves' are, since nothing has the monadic property 'moves', false, and that they should be replaced with the closest truths in the vicinity: 'x moves relative to frame of reference F'. This is, in effect, a move from a judgment of the form:

'X is P'

to a judgment of the form

'xRy'

Here, Kusch observes that the replacing proposition is *not*, as he puts it, 'constructed around' the original 'x is P' (2010, 168). That is, in the case of physical relativism, the relativist is not asking us to replace 'x moves' with

'According to some frame of reference, x moves'.

Rather, the replacing proposition (taking the logical form xRy) is of the form:

'x moves-relative-to(-frame-of-reference)-F1'.

And this is to replace the original predicate ‘moves’ that expresses a monadic property with a new relational predicate, which expresses a dyadic property.

So why, then, does Boghossian think that the *epistemic* relativist – a philosophical relativist – is to be regarded as asking us to move from judgments of the form:

x is P

to a judgment of the form

(X is P) bears [relation] R to S? (2006b, 30)

The relevant difference between the logical forms of the replacement that’s *available* in the two cases can be understood against a background constraint on what is needed to prevent the ‘replacement’ from collapsing into *eliminativism*. If two predicates are no more related than oxygen and phlogiston, then *replacing* one with the other can’t be to *relativise* the first. The first is simply lost. Kusch asks rhetorically: ‘Why is the relation between “moves” and “moves-relative to F1” not like the relationship between “phlogiston” and “oxygen”?’ Here, Boghossian’s own answer is interesting, and it’s one he thinks applies to motion, but not in the case of justification. Specifically, Boghossian (2006, 32) writes that there is a:

more general concept, MOTION, *itself neither absolutist nor relativist*, such that both the absolutist and the relativistic notions could be seen as subspecies of it.<sup>14</sup> [my italics]

And this is why the discovery that the monadic predicate ‘motion’ should be replaced with the dyadic predicate ‘x moves-relative-to (-frame-of-reference)-F1’ is not eliminativism. This rationale betrays an even *more* general principle operating in the background. Boghossian (2006b, 32) even names it: the *requirement of intimacy*:

*Boghossian’s Requirement of Intimacy*: given that a relativist view of a given domain always involves the replacement of the original absolute judgments by certain relational judgments, we need to be shown that these two sets of judgments are sufficiently intimately related to each other, in the sense just gestured at [between motion and relative

motion] to justify our saying that what we have on our hands is relativism and not eliminativism.

But why not think that the epistemic relativist can play the same card, to borrow the logical structure of replacement relativism, as a semantic strategy, in *the unobjectionable physical case*, without collapsing into eliminativism?

The second working part in Kusch's argument is to suggest that such a move is in fact available, because there actually is a general concept JUSTIFICATION that is itself neither relativist nor absolutist. This claim functions importantly in Kusch's overarching modus tollens – viz, that epistemic replacement (as a semantic strategy) relativism is objectionable only to the extent that it is also objectionable in the physical case (where it clearly *isn't* objectionable). In order, though, to liken JUSTIFICATION with MOTION (in so far as both would satisfy Boghossian's criterion of intimacy), Kusch needs to be able to defend the following 'intimacy analogy':

*Intimacy Analogy:* Just as there is there is a more general concept, MOTION, itself neither absolutist nor relativist, such that both the absolutist and the relativistic notions could be seen as subspecies of it, there is a more general concept, JUSTIFICATION, *itself neither absolutist nor relativist*, such that both the absolutist and the relativistic notions could be seen as subspecies of it.

Kusch's defence of this *intimacy analogy* is an innovative one, and I think one of the more important contributions to the recent epistemic relativist literature; and so I'll consider and evaluate it in some detail. Kusch begins by drawing the distinction between three viewpoints: (i) the viewpoint of the *non-philosophical ordinary person*; (ii) the viewpoint of the *absolutist*; and (iii) the viewpoint of the *relativist*. Kusch's non-philosophical ordinary person does not have any philosophical commitments to *meta-epistemic* views (such as absolutism or relativism). Call this the *ordinary person thesis*:

*Ordinary person thesis:* The ordinary person has no commitments to absolutism or relativism (2010, 169).

Kusch supports the ordinary-person thesis by stipulating that the ordinary person is not a philosopher and by supposing that (more or less)

only philosophers have commitments to meta-epistemic views.<sup>15</sup> Kusch combines the ordinary-person thesis with what we can call the *epistemic community thesis*:

*Epistemic community thesis*: The ordinary person, relativist and absolutist can belong to one and the same epistemic community: they can share the same (first-order) epistemic system<sup>16</sup> (2010, 169).

A corollary of the epistemic community thesis is that what distinguishes the ordinary person, the relativist and the absolutist is *second-order*, not first-order commitments (e.g., acceptance/rejection of judgments of the form 'Otto's belief in ghosts is unjustified'). Rather, the relativist and absolutist can be understood as giving the first-order judgment a *different kind of meta-epistemic gloss, even though* (via the epistemic community thesis) the ordinary person, the relativist and absolutist might all share the same (first-order) epistemic system. And it is with this point in mind that we can appreciate how Kusch's view constitutes an – albeit, very different – version of *incompleteness*-theoretic replacement relativism.

As Kusch (2010, 171–72) sums up the thrust of his move:

the particular epistemic judgements and general epistemic principles of the ordinary person are – in the eyes of my relativist – incomplete insofar as they do not express the thought that ours is just one of many equally valid epistemic systems. This incompleteness is not like 'Tom is taller than ...'. In our case what is needed to effect the completion is *the addition of a specific meta-epistemic philosophical gloss*. However, the absence of this specific complement does not leave behind a meaningless torso of words or concepts: it leaves behind the very principle to which the relativist – insofar as he too has been an ordinary person all along – has been, and continues to be committed.<sup>17</sup>

Kusch is thus distinguishing between *first-order incompleteness* and *second-order incompleteness*. Whereas no single epistemic system is singled out in cases of first-order incompleteness (e.g. 'Tom is taller than ...') second-order incompleteness entails just that 'no meta-epistemic stance (of relativism or absolutism) *vis-à-vis* one and the same epistemic system has been determined. Boghossian has collapsed the two forms of incompleteness into one' (2010, 172).

### 6.6.2 Four problems solved?

Kusch's second-order incompleteness-theoretical replacement relativism has the resources to pretty straightforwardly deal with the problems that looked insurmountable to the first-order version of the view Boghossian attacked. As Kusch notes, the *endorsement problem* and the *infinite regress problem* have (on the second-order view) an easy and interconnected solution. Firstly, the endorsement problem; distinguishing between first- and second- order incompleteness undercuts the question: 'how can we endorse incomplete propositions?' This was after all a problem that had bite because, in cases of first-order incompleteness (e.g. Tom is taller than ...), no epistemic system has been singled out. But Kusch does *not* regard such propositions to be first-order incomplete (e.g. p. 172). So the problem doesn't arise. Relatedly, the infinite regress problem does not arise, either, when the relevant incompleteness is regarded as at only the second order. As Kusch puts it, 'the relativist does not need to go through all levels of meta-epistemic complementing before his epistemic system has a content' (2010, 172).

Unsurprisingly, the entailment problem also looks easy to solve now. At least, we no longer have to answer the question of how first-order incomplete propositions can entail one another. Inferential relations are preserved between second-order incomplete statements. Finally, the normativity problem, one that Boghossian took to face *any* form of replacement relativism: Kusch regards this problem as 'real', however, he suggests that at least one viable way to go is to combine relativism with a form of emotivism – *quasi-absolutism* – which Kusch notes has already been defended by, among others, Wong (1984; 2006) and Harman (1996) to get around this sort of problem.<sup>18</sup> At the end of the day, then, Kusch regards replacement relativism as an unproblematic semantic strategy for the epistemic relativist.

#### *Evaluating Kusch's manoeuvre*

It's obvious that Kusch's attempt to resurrect a version of incompleteness-theoretical replacement relativism is much, much better than the version Boghossian challenged. I want to say some positive things about this approach and then raise some challenges.

First some optimism. Kusch is of course right that some cases of *physical* relativism are (obviously) entirely unobjectionable, and further that there is equally nothing intractably unworkable about replacement relativism as a semantic strategy for representing the truths discovered

by the physical relativist. And given that Boghossian's argument is that replacement relativism, as a semantic strategy for representing epistemic relativism, is intractably unworkable, it is very natural to suppose that a promising move for combating Boghossian's argument *strategy* is to attempt to close the gap between the kind of replacement relativism Boghossian finds objectionable in the epistemic case and the kind that is regarded by most everyone as unobjectionable in the physical case.

A second note of optimism: *if* Kusch's *ordinary person thesis* and *epistemic community thesis* are true, then the path does seem open for Kusch to successfully close the gap between physical and epistemic relativism, by showing how, without devolving into eliminativism, the relativist can replace the predicate 'justified' with a long, hyphenated predicate that includes the relativist's preferred meta-epistemic gloss, *viz.*, justified-according-to-epistemic-system-ES1-which-is-one-of-many-equally-valid-epistemic-systems. (Compare: *x* moves-relative-to (-frame-of-reference)-F1). After all, if the ordinary person and epistemic community theses are true, then it looks like Boghossian's requirement of intimacy is going to be met in broadly the same way as Boghossian suggests it is met in the physical case of motion. As Kusch (2010, 173) puts it:

The relativist's relativised concepts are intimately related to the concepts of the ordinary person. The intimate relation is due to the fact that the relativist's successor-concepts preserve the original inferential relations between the ordinary person's concepts.

Provided that Kusch is right that something like a 'quasi-absolutist' strategy is viable for solving Boghossian's normativity problem, I think we should grant Kusch the following conditional: that *if* Kusch's ordinary person thesis and epistemic community thesis are true, then there seems to be no insuperable barrier (of the sort Boghossian took there to be)<sup>19</sup> to replacement relativism as a semantic strategy for modelling the epistemic relativist's core insight. And granting even this conditional should be by Boghossian's lights a serious concession: for one thing, granting Kusch his conditional is enough to cast doubt that we have *a priori* grounds for finding replacement relativist's semantic strategy as off the table for the epistemic relativist. The ordinary person thesis and the epistemic community thesis are, after all, largely beholden to facts which could have been different.<sup>20</sup>

The other side of the coin is this: if the ordinary person thesis is *false*, then it's hard to see how the epistemic community thesis would be true.

After all, if the ordinary person's endorsement of first-order claims *does* involve meta-epistemic commitments – be they relativist or absolutist – then on the very plausible assumption that it's not true that *if ordinary persons had meta-epistemic commitments they would all be the same*, we have no good reason to think that what the epistemic community thesis says is true.

Now, getting to the crux of the matter: is the ordinary person thesis true or false? Kusch – himself aware that he's gained traction against Boghossian even if the ordinary person thesis is not true – offers to, as he says, 'at least briefly explain my intuition that it [the ordinary person thesis] must be roughly on the right track' (2010, 170).

To this end, Kusch offers a defence and an offence. He begins on defence, remarking that: 'Philosophers often count the frequency with which the ordinary person makes non-relativised statements like (1) as conclusive evidence for his commitment to absolutism. I am not convinced' (2010, 170). For one thing, the ordinary person is willing to gloss, for instance, 'Otto's belief in ghosts is unjustified' as 'According to our epistemic system, Otto's belief in ghosts is unjustified.' Secondly, philosophers often 'lament about the allegedly flat-footed relativism of their undergraduates' – undergraduates whom, I take it, Kusch is classing (plausibly) as ordinary untrained people. On offence, Kusch (2010, 170) writes that:

[...] according to my own experience of epistemic discussions with untrained students, when pressed on their stance vis-à-vis the relativism-absolutism opposition, they find it hard to come up with a straightforward answer. This does not of course suggest that philosophically untrained people are epistemic relativists; what it does indicate instead is that being introduced to, and becoming competent in, the practice of epistemic discourse does not involve deciding between epistemic absolutism and relativism. Most of our epistemic discourse functions in ways that do not bring this meta-epistemic alternative into view. And hence ordinary persons tend not to be committed either way.

I am sympathetic to Kusch's remarks here. At least, I think that even if the ordinary person would, contrary to what Kusch is suggesting in the above passage, affirm a meta-epistemic position (be it absolutist or relativist) if presented with a *framed question*, this is still compatible with Kusch's core insight, which is that being competent, as the ordinary person is, in the *practice of epistemic discourse* does not obviously involve any disposition to affirm any particular meta-epistemic position.

Importantly, though, the ordinary person thesis has an ambiguity at its heart. As was discussed in Chapter 1, there are two very different *kinds* of meta-epistemic commitments: articulated meta-epistemic commitments, which one undertakes *via* assertion (e.g. as when Schafer-Landau asserts that moral realism is true), and *revealed* meta-epistemic commitments, which one incurs not by asserting anything but by taking certain things for granted. On the line I advanced in Chapter 1, revealed meta-epistemic commitments are what is pragmatically presupposed by interlocutors engaged in first-order discourse – *viz.* what we find in the *common ground*. And, with reference to first-order debates *in mainstream epistemology*, typical debates (such as, I argued, the debate between Moore and the sceptic, and between Goldman and Feldman about epistemic internalism and externalism) are ones where both sides to the disputes are disposed to behave, in their use of language, *as if* they believe there are epistemic facts with an objective profile. This was tantamount to taking for granted, I argued, a minimal kind of metaepistemological realism.

Granted, the mainstream epistemologist (to whom the argument from Chapter 1 applied) is not the ordinary person. But the distinction between articulated and revealed metaepistemic commitments, as apposite to mainstream epistemology, can easily be transposed to the case of Kusch's ordinary person. For example, suppose two ordinary individuals are arguing about whether Otto's belief in ghosts is justified. Let's follow Kusch and use first-year undergraduates as our interlocutors. Even if we grant Kusch that his undergraduates might well balk if asked whether they regard relativism or absolutism as a correct meta-epistemic position (let's just assume this – I think Kusch is probably right), I suspect that these same undergraduates will, in the conversational moves they make, pragmatically presuppose that there is some non-relative answer to the question of whether Otto's belief in ghosts is justified, and in doing so, take for granted that there is some non-relative fact of the matter. Such a pragmatic presupposition would be evidenced by their (for instance) challenging each other in a way broadly *analogous* to what we find with Goldman and Feldman, Moore and the sceptic, in Chapter 1.

I submit then that, in light of the distinction between articulated and revealed meta-epistemic commitments, we can likewise distinguish two corresponding versions of Kusch's ordinary person thesis:

*Ordinary person thesis-(A)*: The ordinary person has no articulated meta-epistemic commitments to absolutism or relativism (2010, 169).

*Ordinary person thesis-(B)*: The ordinary person has no revealed meta-epistemic commitments to absolutism or relativism.

If the plausibility of the epistemic community thesis depended only on the truth of ordinary person thesis-(A), then Kusch is well positioned to argue for the view that there is a general concept, JUSTIFICATION, of which relativist and absolutist justification are intimately related species. After all, the ordinary person thesis-(A) is plausibly true. But ordinary person thesis-(B) is not true. And why not think that the truth of the epistemic community thesis doesn't *also* depend on the truth of ordinary person thesis-(B)? Alternatively: what principled reason is there for thinking that articulated, but not revealed, meta-epistemic commitments are the only ones that, in the sense relevant to whether Boghossian's intimacy requirement is met in the case of epistemic justification, *count*? Kusch might well have a good answer to this, though I am not optimistic.

## 6.7 Wright versus Boghossian/Kusch: critical discussion

Let's assume for the sake of argument that Kusch *could* satisfactorily address the worry I raised in the previous section for his version of incompleteness-theoretical replacement relativism. In fact, from here on out, we can just hold this assumption in place: that is, let's assume that Kusch's second-order brand of incompleteness-theoretical replacement relativism *can* satisfactorily escape what were, by Boghossian's lights, intractable problems which face any coherent formulation of epistemic relativism on the 'replacement' model.

With this assumption in place I now want to highlight a commitment that is *shared* by Kusch and Boghossian, even though Kusch has, for his part, been defending a view that Boghossian thinks is of an indefensible *type* (e.g. epistemic relativism, formulated on the replacement model). Here, in short, is what Boghossian and Kusch can, despite their differences, agree about:

*Boghossian & Kusch (BK)*: The replacement model offers a satisfactory way to semantically represent what the epistemic relativist wants to say about first-order epistemic judgments.

Boghossian thinks BK is true *even though* he thinks that the replacement model fails to provide a coherent way to formulate the epistemic relativist's thesis. Boghossian after all thinks that epistemic relativism can't

be coherently formulated. Kusch by contrast reveals his commitment to BK because Kusch defends as workable a version of replacement relativism as what he regards as the most defensible version of the thesis and which, *contra* Boghossian, he thinks is not incoherent.

If it turns out that the replacement model is, in fact, not a satisfactory way to semantically represent what the epistemic relativist wants to say about first-order epistemic judgments, then this is a problem for both Boghossian and Kusch. I want to now turn to some considerations Wright (2008, 382–384) brings to bear against Boghossian and Kusch's shared commitment to BK, and then to show how the truth of the Principle of Translation (POT) is relevant to whether we should side with Wright or Boghossian and Kusch.

To make this as clear as possible, let's first consider an objection Wright presents to BK, specifically in the context of challenging Boghossian's formulation of epistemic relativism along replacement lines. And then we can consider how the challenge effectively extends to Kusch's defence of BK in light of the details of Kusch's own proposal.

After quoting Boghossian (2006a, 84) at length, in the passage where Boghossian outlines the three central clauses of epistemic relativism as he sees it (epistemic non-absolutism, epistemic relationism and epistemic pluralism), Wright immediately remarks:

We can envision an epistemic relativist feeling very distant from this characterisation and of its implicit perception of the situation (2008, 383).

And Wright has in mind here especially the epistemic *relationist* leg of the view. This, for ease of reference, was the following:

*Epistemic relationism*: If a person, *S*'s, epistemic judgments are to have any prospect of being true, we must not construe his claims of the form "E justifies belief *B*" as expressing the claim *E justifies belief B* but rather as expressing the claim: *According to the epistemic system C, that I, S accept, information E justifies belief B* (Boghossian 2006a, 84).

But why should the relativist feel distant to this? In one sense, it seems like epistemic relationism is really just a natural extension of other things that the epistemic relativist wants to say in virtue of being an epistemic relativist. That is, epistemic relationism seems to follow naturally from (i) non-absolutism (which epistemic relativists are trivially

committed to) and (ii) what we might call the *vindication* thesis: that, even though first-order epistemic judgments don't aspire to absolute truth, they nonetheless aspire to truth (which is why epistemic relativists talk as though some first-order judgments are true). An explicit commitment to non-absolutism in conjunction with a betrayed acceptance of the vindication thesis would seem to point directly to relationism: that the closest truths in the vicinity of the absolute truths which epistemic relativists must reject are explicitly relational truths which they can accept.

But this passage to epistemic relationism is not as direct as it seems. In order to complete the passage from non-absolutism and the vindication thesis to epistemic relationism, one first needs a reason to think that it is *absolute* truths which are the only truths the relativist can accept, provided the relativist wants to talk as though first-order claims are true while embracing epistemic non-absolutism. Of course, one might think that trafficking in absolute truths about epistemic justification is precisely what the epistemic relativist wants to avoid when purporting to traffic in relative truths about epistemic justification. Though trafficking in absolute truths about epistemic justification is in fact *built into* the relativist's would-be trafficking in relative truths about epistemic justification provided a certain background principle is granted: the very principle Burnyeat (1976) thought we needed in order to make sense of Protagoras – the *Principle of Translation*.

*Principle of translation (POT):* A proposition of the form 'x is F' is true (relatively) for person (a), if and only if, 'x is F for a' is true (absolutely).

Granted, it might be true that even if POT (or, more accurately, a suitably similar principle)<sup>21</sup> is in fact needed to make sense of global relativism, it's not needed to make sense of epistemic relativism. (The two views might after all require different lifelines). But regardless, a background commitment to POT looks like it would neatly explain why one would be inclined to, in characterising the epistemic relativist's take on first-order epistemic judgments, move from epistemic non-absolutism (with a collateral recognition that epistemic relativists talk as though some judgments are true) to epistemic relationism.

Though Wright doesn't speak of the principle of translation specifically (i.e. by referring explicitly to Burnyeat's principle), he certainly, in the following passage, intimates it when suggesting that we're not

giving the epistemic relativist a fair shake when attributing to her the relationist clause. Here's Wright:

Boghossian is taking it that the truth of unrelativised claims such as (1), straightforwardly construed, is intelligible only if supported by absolute facts about justification. Since – for the relativist – there are no such facts, sentences like (1), if they “are to have any prospect of being true”, have to be construed as making some other kind of claim, whose truth can be supported by facts of a kind that relativism can consistently countenance; and then the only salient candidate facts are the explicitly relational ones invoked by the epistemic relationist clause, B (2008, 383).

And then, Wright adds this critical assessment:

But this is simply tantamount to the insistence that if relativism is to be in any position to regard claims such as (1) as true at all, then it is obliged so to construe their content as to enable them to be made *absolutely true* or false by some class of facts that it countenances. That is just to fail to take seriously the thesis that claims such as (1) can indeed be true or false, albeit, *only relatively so*.<sup>22</sup> (*Ibid.*, 383)

Wright for his part thinks that there is a more charitable formulation of epistemic relativism, one where epistemic relativism is situated within what he calls the ‘new age’ framework – *viz.*, an assessment-sensitive semantic framework, *a la* MacFarlane (2005; 2009; 2014). We'll explore this proposal in depth in the next chapter. But for now consider where things stand for both Boghossian and Kusch. If Wright is right in his critical remarks, then Boghossian has betrayed himself as committed (in moving from epistemic non-absolutism to epistemic relationism) to something like the principle of translation – in a way that does not take seriously that (1) – claims can indeed be true or false, albeit only relatively so. And this is one respect in which Boghossian's commitment to BK is problematic by Wright's lights.

As for Kusch: I've suggested already that he is also tacitly signing up to BK. By positively defending the viability of a version of epistemic replacement relativism, at least, Kusch is regarding the replacement strategy as a satisfactory way to formulate the view. (And that's enough to be committed to BK, which Wright rejects). After all, on Kusch's own formulation, the meta-epistemic gloss that the epistemic relativist

gives first-order epistemic judgments is explicitly relational as well. The logical form is different, granted. For Kusch, when the first-order epistemic judgment is completed with a meta-epistemic gloss by the relativist, we have a proposition that replaces a monadic predicate with an explicitly relational dyadic predicate, and then includes the further the qualifier that the judge's epistemic system is just one of many equally valid epistemic systems. This no more by Wright's lights preserves the insight that what the relativist wants to say about first-order judgments is that they can be true or false, albeit only relatively so, than does Boghossian's own preferred formulation of the relationist clause. In short, both Boghossian and Kusch read the relativist as replacing one truth with another, whether the replaced truth is one that 'smuggles' the relativisation into the content of the claim that's regarded as true in such a way that the explicitly relational truth could be absolutely true. If *this* is not to take the relativist seriously, then neither Boghossian nor Kusch has taken the relativist seriously.

But what of the antecedent here? *Is that* not to take the relativist seriously? That's a substantive question. Kusch might well say that it is to take the relativist seriously, because after all he's an epistemic relativist and he is not put off by relationism as an element of the view. And indeed, we can envision a stronger claim here: Burnyeat thought that (at least, in the case of Protagoras) we aren't taking him seriously *unless* (not: *a la* Wright, *if*) we are prepared to attribute a background commitment to the principle of translation: according to which a commitment to a given relative truth just is a commitment to an explicitly relational absolute truth.

The principle of translation is thus a knife that might be wielded in both ways. Depending on whether one regards keeping this principle in play as *essential to* or *incompatible with* taking the epistemic relativist seriously, one might either view BK as problematic (as Wright does, and as we'll see in the next chapter, MacFarlane does) or not (as Boghossian, Kusch and Burnyeat do). Attempting to answer the philosophical question of what it is to take epistemic relativism seriously looks like a trap, and I'm not going to try to adjudicate this by reflecting on what it is to take any view seriously, etc. Rather, there might well be a way to move forward and potentially settle the issue in favour of Wright (in a way that does not beg the question either way about what is needed to take epistemic relativism seriously).

On this point, I want to turn to an objection Wright levels against replacement relativism (replete with its relationist clause) as such, which might well have the effect of raising an independent reason to be

suspicious of BK, one that does not involve any posturing on the matter of what epistemic relativists *really* stand for.

As Wright (2008, 383) sees it:

[...] any attempt to construe claims of type (1) in accordance with the broader strategy proposed in clause B – i.e., as in effect elliptical expressions of certain neighbouring explicitly relational claims, is completely unworkable...it involves a confusion between making a judgment *in the light of* certain standards and judging that those standards mandate that judgment. A judgment of the latter kind may be perfectly rationally endorsed by a thinker who is not at all inclined to the original judgment – because he does not share the standards in question.

Here we need to be careful. Wright's problem – call it the *standards-conflation problem* – looks very much like another *aspect* of the normativity problem. The normativity problem, recall, is that 1-style judgments are normative, but their explicitly relational replacements aren't. Now, Boghossian, in his argument against the coherence of epistemic replacement relativism, raised the normativity problem *himself* (2006b, 25, 2006a, 87). So this point in the dialectic might seem a bit confusing. How can Wright's arguing that the standards-conflation problem reveals replacement relativism to be unworkable count against BK (which Boghossian accepts) given that Boghossian agrees completely with Wright on the general thrust of the point – *viz.*, that the replacement relativist's relationist clause fails to appropriately preserve the *sense* in which one is accepting, or normatively endorsing, the first-order (type 1) epistemic judgment, when one puts that judgment forward?

The answer is simple and subtle. Wright thinks that epistemic relativism actually *can* in fact be formulated in a way that's doesn't straightforwardly generate this problem, namely, along the lines of the 'new-age' template. On the new-age template, which drops the relationist clause, the relativist doesn't encounter the normativity problem; as the relevant 1-type proposition is made *in light of* certain standards, it is not replaced by a judgment that some standards mandate a particular judgment. We'll engage with the more technical details of this kind of framework in the next chapter, but we've already now got an answer to why Boghossian and Wright's agreement with the thrust of the normativity problem is compatible with Wright using this in a way that counts against BK. The reason is that Wright is taking the problem as evidence that there's a *better* way to formulate the view, whereas Boghossian simply

took it as evidence that the epistemic relativism cannot be coherently formulated.

Kusch might well have some cause to object here. Recall that Kusch advanced a 'quasi-absolutist' strategy for making sense of the normativity of the relevant first-order judgments, even though they are on his model articulated by the relativist in a way that is explicitly relational (*via* the dyadic predicate). The key to Kusch's quasi-absolutist move was to (*a la* Wong and Harman) exploit the way we sometimes use epistemic terminology to *express* our approval of certain epistemic systems. For example: 'Otto's belief in ghosts is **unjustified**' where the emboldened use indicates approval. Kusch thinks this is enough to preserve the normativity of the first order judgments (2010, 174). As he puts it:

A sentence is normative if it expresses a norm. And being a norm is related to the phenomena of praise and criticism, approval and disapproval... When I utter [sic. Otto's belief in ghosts is **unjustified**] I criticise Otto's belief in ghosts; when you utter [sic. Otto's belief in ghosts is **justified**] you refuse to criticise him (perhaps indicating criticism of those who would criticise him).... My criticism of Otto's belief involves my disapproval of epistemic systems that permit (and require) him to believe in ghosts. Your regarding his belief as justified you to approving epistemic systems that assess his belief positively. (*Ibid.*, 174)

Kusch also thinks this move helps to preserve disagreement, but never mind this for now. What's relevant here is that Kusch's defence of his version of replacement relativism against the normativity objection reveals a subtle way that the normativity objection and Wright's standards-conflation problem come apart, *even if* the standards-conflation problem is (as I've already suggested) closely connected to the normativity problem. To appreciate this point, consider that the rationale Wright gives with respect to the standards-conflation problem (in the service of suggesting that attributing to the relativist the relationist clause is to not take her seriously) is one that will continue to support the view that Kusch hasn't given the relativist's position a fair shake *even if* we grant that Kusch's quasi-absolutism suffices to diffuse the normativity problem.

This line of thinking takes as a starting point the observation that Kusch's quasi-absolutist strategy is one that is *built into* his relationism. And it was the relationism itself that falls prey to the standards-conflation problem. So even if Kusch can (*via* the quasi-absolutist

strategy) preserve the normativity of first-order judgments, Wright will remain in a position to insist that Kusch's replaced judgments are, and there's no escaping this point – that is, when the first-order judgments are second-order completed with the relativist's meta-epistemic gloss – still *judgments about standards mandating a judgment*. And this is true, again, even if we concede to Kusch that the quasi-absolutist move is permissible to preserve normativity.

Though, in the face of this line, Kusch might yet have a card (with two sides) to play. Consider this reply by Kusch, which I envision as having two parts. Here's part one: 'provided my version of replacement relativism can satisfactorily address the normativity and the endorsement problems, then how serious can the standards-conflation problem really be, as a strike against the view?' This first part of the envisioned answer is a *mitigating* reply: it is meant to soften the blow of the charge. The second envisioned response is stronger. According to the second leg of the reply, the standards-conflation problem is granted as a real problem, but Kusch can then suggest he can solve it despite appearing to be unable to. This leg of the reply takes some delicacy. For clarity sake, let's look at a concrete example. Kusch's relativist is reconstructing 'Otto's belief in ghosts is unjustified' as expressing this proposition:

*Kusch Replacement Proposition (KRP):* Otto's belief in ghosts is unjustified-according-to-epistemic-system-ES1-which-is-one-of-many-equally-valid-epistemic-systems (2010, 171).

Is KRP an example of judging in light of standards or judging that standards mandate a judgment? Granted, it *looks* like the latter, doesn't it? But maybe this is deceptive. Again, helping himself to the quasi-absolutist strategy, Kusch can begin by pointing out that his relativist retains *approval*. That is, that when the relativist advances KRP (for instance, in a disagreement with another relativist who substitutes ES1 with ES2 (e.g. 2010, 174–175), the first relativist can be understood as approving a different epistemic system than is the relativist's (also relativist) interlocutor (*ibid.*, 175). If Kusch has secured this point, then it seems he has a potentially interesting reply waiting in the wings to the standards-conflation problem: Wright accuses the replacement relativist (of any sort) of confusing (i) judging in light of standards; and (ii) judging that standards mandate a judgment. But now Kusch might be in a position to insist, with reference to his position, that this is a false choice – that these are not *mutually exclusive*. On the envisioned 'false choice' reply, Kusch can suggest that *even if* KRP is an instance of judging

that standards mandate a judgment, it needn't follow that it's *not also* instance of judging in light of standards – something Kusch can make a case for by appealing to his quasi-absolutist move which (suppose) we grant allows him to preserve that the relativist can retain 'approval'.

There might be a counterreply, though. Maybe Wright, in issuing the standards-conflation objection to relationist views, wants to say something even *stronger*: maybe the *real* thrust of the standards-conflation problem is supposed to be that, in order to interpret the relativist charitably, we should want the replacement to *positively not* involve a judgment that standards mandate a judgment, *regardless* of whether the judgment in question can also be understood as being made *in light of* a standard. On this stronger reading of the charge, Kusch succumbs to the standards-conflation problem even if he retains that his view preserves, in the relevant sense, judging in light of a standard. But adjudicating the merits of this counterreply is difficult without taking a stand on what it is the epistemic relativist is *supposed to stand for*, and here we risk posturing.

## 6.8 Concluding remarks

So where does this leave us? I want to conclude by highlighting three points, meant to connect where we stand now with where we're going in the final three chapters. Firstly, properly adjudicating the efficacy of Wright's critique of Boghossian and Kusch is, in light of what's been said so far, premature. After all, the registering of a verdict here – on the matter of whether Wright is correct that Boghossian and Kusch are ultimately presenting the relativist's idea in less than the most charitable way – cannot be done in an appropriately informed way without substantively engaging with the matter of whether the 'assessment sensitive' new age model in fact does provide, as Wright thinks, an all-things-considered better alternative model for capturing the relativist's core insight.

Secondly, assessing the viability of epistemic relativism on the new-age model, however, is – and this is important – not *just* going to be a matter of plugging epistemic relativism into a MacFarlane-style assessment-sensitive framework and seeing whether we can fend off the stock objections – *viz.*, the normativity problem, endorsement problem, entailment problem, etc. – to the replacement model in a more satisfactory way than the replacement model itself can be defended against these objections.<sup>23</sup> This is because, and this point bears emphasis: the *sources of philosophical motivation* that drive MacFarlane to give 'knows' a relativist semantics

are in fact very different from the sources of philosophical motivation that have underwritten the kinds of arguments for epistemic relativism we've been considering thus far – e.g. in chapters 3–5. As MacFarlane sees it, a *truth-relativist* semantics better explains our practices of attributing knowledge than competing semantics for 'knows', such as classical invariantist, subject-sensitive invariantist, contextualist semantics. For MacFarlane, the philosophical motivations for a truth-relativist treatment of knowledge ascriptions are primarily to do with language and semantics. An evaluation of the viability of epistemic relativism on a truth-relativist framework will involve at least, then, an evaluation of how a MacFarlane-style semantics for 'knows' fares against other competing semantic theories. In the context of this evaluation, it will also be important to assess the implications a truth-relativist semantics for 'knows' has more broadly in epistemology, and indeed, how plausible this ultimately is, a matter about which MacFarlane himself hasn't said much.

Finally, a big-picture remark: note that while I've withheld judgment (until we engage with MacFarlane in the final chapters) on the matter of whether, *if philosophical arguments for epistemic relativism are compelling*, the view is best formulated within a replacement semantic template, or within a MacFarlane-style semantic template, I have *not* withheld judgment on a philosophical question that was posed toward the end of Chapter 1. I want to conclude this chapter by briefly revisiting that question and where things now stand with respect to it.

The simple version of the question was whether epistemic relativism should be taken more seriously by mainstream epistemologists than it in fact is taken – *viz.*, most first-order debates in mainstream epistemology, I've suggested in Chapter 1, *proceed* exactly as if epistemic relativism is false. *Is this a mistake?* Or, by contrast, is it no great loss that the kind of meta-epistemology I argued to be pragmatically presupposed by projects in mainstream epistemology *seems* to entail the falsity of epistemic relativism? *This* depended, I suggested (though there are more details at the end of Chapter 1 I won't revisit here) importantly on whether the relativist gives us a good reason to accept her picture of epistemic facthood. And *this*, depended on the matter of whether *there are in fact compelling reasons to accept the picture of the landscape the epistemic relativist is selling*.

The case developed over Chapters 3–5 indicates that mainstream epistemology has suffered no great loss for proceeding as though epistemic relativism is false; at least, chapters 3–5 suggested that *familiar* argument strategies for epistemic relativism have ultimately failed to motivate

relativism over other alternatives, particularly in many cases, scepticism, and so have provided no compelling reason to embrace the epistemic relativist's picture of epistemic facthood. That MacFarlane's motivation for a brand of 'new age' epistemic relativism is on an entirely different kind of philosophical basis than the more traditional kinds of rationales for epistemic relativism surveyed over Chapters 3–5 is thus very important. We have *prima facie* reason to think that the *kinds* of familiar problems – especially, the problem of failing to motivate relativism over scepticism – which plagued more traditional accounts, might be entirely orthogonal to MacFarlane's version of the view. After all, the matter of whether a truth-relativist semantics outperforms contextualist or invariantist semantics for knowledge attributions in no discernible way gives rise to the old dialectical worry – that if you follow the sceptic too far, it will be difficult to part ways.<sup>24</sup> If a MacFarlane-style semantics for 'knows' is plausible, then we can anticipate that it won't be able to be dismissed on the basis of the same *kinds of considerations* that have been appealed to in the service of combating philosophical motivations for relativism that have been explored thus far.

# 7

## A Different Kind of Epistemic Relativism

*Abstract.* If John MacFarlane (2014) is right about what makes a philosophical position relativistic in a philosophically interesting sense, then many of the views under the description of epistemic relativism discussed so far fall on the uninteresting side of the line. This chapter has three main goals. Firstly, I outline what makes a view count as interestingly relativistic for MacFarlane, namely, that what he calls a *context of assessment* is afforded a significant semantic role. Next, I detail MacFarlane's rationale for being a relativist in this sense about propositional *knowledge*, by defending the view that claims of the form 'S knows that *p*' should be understood as getting a truth value only relative to a context of assessment. Though this takes a bit of stage setting – some positioning within contemporary work in the philosophy of language – the majority of the chapter thereafter will be critical. MacFarlane's (2014) presentation of his proposed truth-relativist semantics for knowledge attributions keys relevant alternatives to the context of assessment. I argue here in detail why this is unworkable. It is then shown that even if we retreat to a more generic formulation of MacFarlane's view – his (2005) presentation which adverts to the more abstract notion of an 'epistemic standard' as what's keyed to the context of assessment – the *kind* of problems formulated against his 2014 view can be recast. It is concluded that an underlying Achilles' heel for MacFarlane's view is its revealed incompatibility with a plausible thesis in epistemological meta-theory called *epistemic anti-individualism*, the thesis that what converts true belief to knowledge can supervene partly on elements of one's local and/or modal environment. In short, if MacFarlane is right that assessment sensitivity is the mark of relativism, it's hard to see how *knowledge* is relative.

## 7.1 Relativism, revisited: where MacFarlane draws the line

As noted in the previous chapter, Wright's (2008) criticism of Boghossian's (2006a) formulation of epistemic relativism was, in short, that Boghossian, in virtue of including the relationist clause in his formulation of the epistemic relativist's view, simply wasn't taking *seriously* the idea that some first-order epistemic judgments might be 'true, albeit, only *relatively* so' (Wright 2008, 383). We'll begin now to explore more carefully what this charge amounts to.<sup>1</sup>

John MacFarlane, in a number of recent works,<sup>2</sup> advances a principled way of thinking about what counts as *philosophically interesting relativism*, one which makes clear where the 'line' is supposed to be between truth *absolutism* on the one hand, and a position that relativises truth in a way that *parts ways* with the absolutist on the other. Wright's charge that Boghossian is not taking seriously what the relativist is saying can be understood as, at bottom, the charge of not attributing to the relativist a kind of view that falls on the 'relativist' side of this line. But to appreciate this point requires that we situate the discussion within analytic philosophy of language.<sup>3</sup>

To begin with, consider a very simple insight, familiar from Kaplan (1989), which is that the truth-value of an utterance can depend on the context in which it is uttered. For example, if Sir David Attenbrough says 'I have been to Antarctica' and David Boreanaz says the same sentence, David Attenbrough's utterance may be true and David Boreanaz's false. In such cases, the context of utterance plays a role in determining which proposition the sentence 'I have been to Antarctica' expresses, e.g., whether it expressed *Sir David Attenbrough has been to Antarctica* or *David Boreanaz has been to Antarctica*.

While this point is mundane in the cases where statements feature overt indexical expressions (e.g. 'I', 'tomorrow', 'here'), some philosophers hold that terms which are not overtly indexical should nonetheless be treated like such. Take, for example, moral terms, e.g., 'wrong'. According to Harman (1996) and Dreier (1990), for example, a statement of the form 'Murder is wrong' is roughly equivalent to 'Murder is wrong according to the moral system I accept'. And so, two utterances of (say) 'Murder is wrong' can – at least on this kind of view – differ in truth-value provided they are uttered by speakers who accept different moral systems.

While *moral contextualists*<sup>4</sup> treat moral terms as indexicals in this way, *epistemic contextualists* give (certain) epistemic terms – typically, ‘knows’ – this kind of semantic treatment, whereby the proposition that is expressed by a sentence attributing knowledge (e.g. in the canonical form, ‘S knows that *p*’) depends upon the context in which ‘S knows that *p*’ is *used*, and so utterances of a knowledge-attributing sentence (e.g. ‘Keith knows that the bank is open’), made in *different* contexts, can vary in truth value.<sup>5</sup> In the simple case of *attributor contextualism*, the attributor’s epistemic standards are what determines what proposition is expressed in a given context of use; in this respect, the truth value of knowledge ascriptions (on attributor contextualism) *depends* on the attributor’s epistemic standards.

*Question:* given that contextualists about knowledge allow that the truth-value of knowledge attributing sentences is (in an obvious sense) *relative* – *viz.*, relative to the context in which the attribution is made – why is it that contextualism is not *itself* simply relativism, or at least a version of it? And even more: given that contextualists allow for the truth-value of knowledge-attributing utterances to depend on *epistemic standards*, contextualists are putting forward a view that bears obvious commonalities with Rorty-style cultural relativism about epistemic justification. But wasn’t Rorty’s view *relativism*?

*Answer:* in one wide, and also fairly intuitive,<sup>6</sup> use of the word ‘relativism’ – one which MacFarlane for his part rejects as uninteresting – the answer is that contextualists simply *are* relativists of a certain stripe. After all, in the epistemic case, they (following the ‘arity’ insight outlined in Chapter 2) surprise us with their purported discovery that ‘knows’ is implicitly relational in a way that we perhaps didn’t expect. And moreover, contextualists do tell us that statements of the form ‘S know that *p*’ do not receive a truth-value unless an epistemic standard is specified, rendering knowledge claims obviously ‘relative to an epistemic standard’ in a sense that will sound pleasing to the ear of some philosophers who fly explicitly the banner of relativism.<sup>7</sup> And so *if being committed to these things is enough to make one an epistemic relativist*, then contextualists are epistemic relativists.

But of course, maybe being committed to these things is in fact *not* enough. After all, according to a contextualist, even if knowledge ‘is relative to an epistemic standard’ in the sense just articulated, it remains that – for the contextualist – a particular occurrence of ‘knows’ used in a particular context, has its extension *absolutely*.<sup>8</sup> A key theme of John MacFarlane’s work over the past decade has been to suggest that simply relativising propositional truth to what *seem like exotic parameters* (e.g.

other than worlds and times – such as judges, perspectives, or standards (including *epistemic* standards) – is not in itself ‘enough to make one a relativist about truth in the most philosophically interesting sense’. The interesting divide, he argues, ‘is between views that allow truth to vary with the context of assessment and those that do not’ (2014, vi).<sup>9</sup> The ‘line’, for MacFarlane, is thus demarcated with reference to his notion of a context of assessment,<sup>10</sup> which he says (2014, 60) we can distinguish as related to the familiar notion of a *context of use* in the following way: whereas a *context of use* is a possible situation in which a sentence might be used and where the agent of the context is the user of the sentence, a *context of assessment* is a possible situation in which a *use* of a sentence might be *assessed*, where the agent of the context is the *assessor* of the use of a sentence (*ibid.*, 60).

On MacFarlane’s view, *truth-relativism* (of the philosophically interesting sort, i.e., worthy of the term ‘relativism’) with respect to utterances in area of discourse D is, in short, the claim that the truth of S’s D-utterance *u* depends (in part) on a *context of assessment*;<sup>11</sup> that is, what S asserts, *u*, gets a truth value – according to the truth-relativist’s D-semantics – only once the D-standard of the *assessor* is specified. And, correlatively, *independent* of the specification of such a standard, S’s *u* assertion simply lacks a truth-value much as, by comparison, indexical expressions such as ‘I have been to Antarctica’ do not get a truth-value independent of contextual facts about the context of *use* (i.e. the context in which the utterance is made).<sup>12</sup>

For MacFarlane, then, the answer to the question of why contextualism is not on the interesting side of the relativist’s line is that while the contextualist can, no less than the relativist, recognize a ‘standards’ parameter, for the contextualist, its value will be supplied by the context of *use*, whereas the relativist (proper) takes it to be supplied completely *independently* of the context of use, by the context of *assessment*.<sup>13</sup>

I’ll be most concerned, in what follows, with the prospects of a kind of *epistemic* relativism which MacFarlane embraces by defending the view that ‘knows’ is assessment-sensitive. But first two more general points will be useful in thinking about assessment-sensitivity as the ‘mark’ of relativism.

The first general point concerns how the *rationale* for embracing a MacFarlane-style relativist semantics for ‘knows’ should be understood as differing from the kind of rationale we find in Lewis’s (1980) and Kaplan’s<sup>14</sup> (1989) foundational work in semantics according to which sentence truth was relativised to familiar parameters such as worlds, times and locations. The important point here is that while Lewis’s and

Kaplan's *reasons* for 'proliferating' parameters were primarily based on considerations to do with intensional operators, the more contemporary reasons (e.g. as appealed to by MacFarlane and other 'new relativists') for adding a standards parameter (i.e. in the context of assessment) are often to do with respecting linguistic use data, e.g. such as disagreement data. For example, those who endorse truth-relativism about predicates of personal taste, (e.g. Lasersohn 2005; Kölbel 2003, MacFarlane 2014) take a truth-relativist semantics to better explain our patterns of using terms like 'tasty' than do competing contextualist, sensitive and insensitive invariantist semantics. Accordingly, *defending* new-age relativism typically involves, for some area of discourse D, a philosophical comparison of costs and benefits of different competing semantic approaches to the relevant D expressions, replete with a case for thinking that the truth-relativist all-things-considered performs the best. A familiar such claimed advantage by a MacFarlane-style truth-relativist is that the kind of 'subjectivity' (e.g. standards-dependence) the contextualist claims the traditional invariantist can't explain can be captured by the relativist without – or so the relativist tells us – 'losing disagreement', where losing disagreement is a stock objection to contextualism in areas where disagreements appear genuine.

The second general point to note is that *if*, as MacFarlane tells us, assessment-sensitivity is the *essence* of philosophically interesting relativism, then any epistemic relativism worth its stripes will be a view on which, for example, Paul's assertion that 'Copernicanism is justified,' can be construed as *true*, to use Wright's language, '*albeit, only relatively so*' namely, only relative to *a context of its being assessed as true or false*. And if the relativist is really *best* understood as thinking of relative truth in this way, then it should be clear why engrafting *epistemic relationism* into a formulation of epistemic relativism would be objectionable. Recall, again, that Boghossian's epistemic relationist clause says:

*Epistemic relationism*: If a person, S's, epistemic judgments are to have any prospect of being true, we must not construe his utterances of the form "E justifies belief B" as expressing the claim *E justifies belief B* but rather as expressing the claim: *According to the epistemic system C, that I, S, accept, information E justifies belief B*. (Boghossian 2006a, 87)

In short, epistemic relationism – and MacFarlane's view that 'The threshold of relative truth is only crossed when we give a semantically significant role to the context of assessment' are mutually exclusive. Pared down to the nub of the matter: the Boghossian-style relationist

clause *insists* that the relativist must make a claim on what MacFarlane regards as the *uninteresting* side of the relativist/absolutist divide.<sup>15</sup> In contrast, MacFarlane *insists* that the relativist, *qua relativist*, not make a claim on that side of the line. This being said, it should be unsurprising why MacFarlane (2014, 33, fn. 5) calls Boghossian's relativist something *other* than a relativist: a contextualist.<sup>16</sup>

The past decade has seen an explosion of proposals where various expressions, in various domains, are claimed to fall on MacFarlane's side of the interesting divide. In particular, 'truth-relativist' semantics (of which MacFarlane's brand of assessment-sensitivity is the most prominent example) have been defended recently for: predicates of personal taste (Lasersohn 2005; Kölbel 2003, MacFarlane 2014), epistemic modals (Egan 2007; Egan, Hawthorne & Weatherson 2007; MacFarlane 2011, Stephenson 2007), future contingents (MacFarlane 2003), indicative conditionals (Weatherson 2009; Kolodny & MacFarlane 2010), gradable adjectives (Richard 2004) and deontic modals (Kolodny & MacFarlane (2010) and MacFarlane (2014, ch. 11)).<sup>17</sup>

Most relevantly for our purposes, MacFarlane<sup>18</sup> claims that *knowledge attributions* are assessment-sensitive, and it's to this claim that I'll now turn – first expositively and then critically.

## 7.2 Relativism about knowledge attributions

In three different places, MacFarlane (2005, 2009, 2014) has argued that *knowledge attributions* of the form 'S knows that *p*' are assessment-sensitive. The focus of his presentation has varied across these three defences of the view, but one core strand of thought resurfaces each time. For ease of convenience, we can call this MacFarlane's 'master' argument for an assessment-sensitive semantics for knowledge attributions.

*Master Argument for Assessment Sensitive Semantics for Knowledge Attributions*

- (1) Standard invariantism, contextualism and SSI all have advantages and weaknesses.<sup>19</sup>
- (2) Relativism preserves the advantages while avoiding the weaknesses.
- (3) Therefore, *prima facie*, we should be relativists about knowledge attributions.

The remainder of this chapter will be concerned with three principal objectives. Firstly, I want to make clear why MacFarlane thinks that

(1) and (2) of the master argument are true, and so why he thinks we should embrace a relativist treatment of 'knows'. In doing so, I'll draw primarily from MacFarlane's most recent presentation of his relativist treatment of 'knows', one which gives the notion of *relevant alternatives* a central place.<sup>20</sup> The second principal objective will be critical one. I'll raise several problems related to MacFarlane's reliance on relevant alternatives in formulating his position; finally, I'll detail a wider kind of problem for his strategy, one which does not depend on his recent 'relevant alternatives' formulation.

### 7.2.1 Problems with invariantism, contextualism and SSI

As MacFarlane sees things, each of the three standard views of the semantics of knowledge-attributions – *standard invariantism*, *contextualism* and *subject-sensitive invariantism* (SSI) – has a grain of truth to it, as well as an 'Achilles heel: a residuum of facts about our use of knowledge attributions that it can explain only with special pleading' (2005, 197). His most recent way of making this point relies on a kind of sceptical 'conundrum', one which arises in light of our ordinary practices of attributing knowledge, and which he uses as a frame of reference for magnifying what he regards as the salient weaknesses of the three standard views.

*MacFarlane's Conundrum:* If you ask me whether I know that I have two dollars in my pocket, I will say that I do. I remember getting two dollar bills this morning as change for my breakfast; I would have stuffed them into my pocket, and I haven't bought anything else since. On the other hand, if you ask me whether I know that my pockets have not been picked in the last few hours, I will say that I do not. Pickpockets are stealthy; one doesn't always notice them. But how can I know that I have two dollars in my pocket if I don't know that my pockets haven't been picked? After all, if my pockets were picked, then I don't have two dollars in my pocket. It is tempting to concede that I don't know that I have two dollars in my pocket. And this capitulation seems harmless enough. All I have to do to gain the knowledge I thought I had is check my pockets. But we can play the same game again. I see the bills I received this morning. They are right there in my pocket. But can I rule out the possibility that they are counterfeits? Surely not. I don't have the special skills that are needed to tell counterfeit from genuine bills. How, then, can I know that I have two dollars in my pocket? After all, if the bills are counterfeit, then I don't have two dollars in my pocket.<sup>21</sup> (MacFarlane 2014, 177)

MacFarlane articulates the form of the conundrum argument as follows:

- (i)  $p$  obviously entails  $q$ . [premise]
- (ii) If  $a$  knows that  $p$ , then  $a$  could come to know that  $q$  without further empirical investigation. [1, Closure]
- (iii)  $a$  does not know that  $q$  and could not come to know that  $q$  without further empirical investigation. [premise]
- (iv) Hence  $a$  does not know that  $p$ . [2, 3, modus tollens]

*Standard (insensitive) invariantism*,<sup>22</sup> the view that the epistemic standards that must be met for 'S knows  $p$ ' to be true are not (in any way) context sensitive,<sup>23</sup> faces two central problems, by MacFarlane's lights. Both problems are familiar. Firstly, standard invariantism has trouble making sense of the *variability* of our willingness to attribute knowledge. Secondly, standard invariantism seems stuck with an unhappy choice of either: embracing *scepticism* (if the invariantist simply accepts (iv)), embracing *dogmatism*<sup>24</sup> (if the invariantist tries to avoid the sceptical conclusion (iv) by rejecting (iii)), or rejecting the *closure principle* which licenses the move from (i) to (ii) – viz., the principle that (as MacFarlane states it): '*if a knows that  $p$ , and  $p$  obviously entails  $q$ , then  $a$  could come to know  $q$  without further empirical investigation*' (2014, 177).

By contrast, *contextualism*<sup>25</sup> offers a way to avoid each of these problems facing standard invariantism.<sup>26</sup> Unlike the invariantist whose position is at tension with data about the variability of our willingness to attribute knowledge, the contextualist has an explanation to offer for this variability: namely, our willingness to attribute knowledge varies across contexts because what is meant by 'knows' is sensitive to the context in which it is used. As MacFarlane writes, 'On the most natural form of this view, 'knowing' that  $p$  requires being able to rule out contextually relevant alternatives to  $p$ . *Which alternatives are relevant depends on the context*'. For instance, and with reference to MacFarlane's Conundrum, when I'm first asked whether I know ( $p$ ) – that I have two dollars in my pocket – 'knowing' that  $p$  requires I need only to be able to rule out very basic alternatives (e.g. that I didn't already spend the \$2); I needn't be able to also rule out that my pockets have been picked to count as 'knowing'<sup>27</sup> (*ibid.*, p. 177). Though when someone asks me whether my pockets have been picked, then 'knowing' requires ruling out this alternative, and if I can't, then the standard required for 'knowing' in this context is not met.

Contextualism can make sense not only of the variability of our willingness to attribute knowledge,<sup>28</sup> but it also avoids the unpalatable dilemma facing standard invariantism: reject closure or embrace scepticism or dogmatism. As the standard line goes, contextualists needn't be tarred as sceptics or dogmatists because they can in fact *preserve* closure, at least, within any one context of use. So contextualism is looking pretty good.

However, although treating 'knows' like 'tall' – where the meaning of knows depends on the context in which it is being used – offers a nice escape route (*vis-à-vis* MacFarlane's Conundrum), there are other respects in which treating 'knows' like 'tall' raises new problems. For example, an apparent disagreement between A and B about whether Michael Jordan is tall quickly is revealed to be no disagreement at all when it is clear to both parties that A means 'tall for a given person' and B means 'tall for an NBA player'. However, as MacFarlane notes, things are different with 'know'. He writes:

If I say "I know that I have two dollars in my pocket," and you later say, "You didn't know that you had two dollars in your pocket, because you couldn't rule out the possibility that the bills were counterfeit," I will naturally take your claim to be a challenge to my own, which I will consider myself obliged either to defend or to withdraw. It does not seem an option for me to say, as the contextualist account would suggest I should: "Yes, you're right, I didn't know. Still, what I said was true, and I stick by it. I only meant that I could rule out the alternatives that were relevant then." Similarly, the skeptic regards herself as disagreeing with ordinary knowledge claims – otherwise skepticism would not be very interesting. But if the contextualist is right, this is just a confusion. (*Ibid.*, p. 181)

And here's where the special pleading comes in. The contextualist can attempt to say that our taking each other to agree/disagree with each other in the relevant kinds of cases is just a mistake of some sort. But this is a double-edged sword – as MacFarlane rightly observes, the more *speaker error* the contextualist must posit to explain the way we use 'knows', the *less* the contextualist can rely on the way we use 'knows' to support contextualism.<sup>29</sup> While contextualism does better than standard invariantism in that it avoids the dilemma raised to standard invariantism, standard invariantism makes better sense of disagreement.<sup>30</sup>

By comparison with insensitive invariantism and contextualism, *subject-sensitive invariantism* might have the best offer to make yet. According to SSI, whether my utterance of 'Dan knows his car is in

the parking lot' is true does depend on context, though in a different sense than it does for the contextualist: rather than to depend on what alternatives *I* (the utterer of the sentence) can rule out (e.g. whether or not I know there are no thieves lurking nearby) what matters on SSI is whether Dan, the subject of the knowledge attribution, can rule out the alternatives relevant to *his* practical environment. More generally, on SSI, the truth of the target claim is sensitive not to the context of *use* but rather to the *circumstance of evaluation* (in the sense of Kaplan (1989)) – in the same sense that, to use an example from MacFarlane (2005, 2–3), 'the distance expressed by "as far apart as Mars and Jupiter" varies with the circumstances (for instance, the time) of evaluation';<sup>31</sup> more specifically, on SSI, "knows" 'invariantly expresses a property whose extension at a circumstance of evaluation depends on features of the subject's practical situation'.<sup>32</sup>

The 'SSlist' can make sense of disagreement, given that 'knows' is not being treated like 'tall' and further, the SSlist unlike the *insensitive* invariantist can make sense of variability in willingness to attribute knowledge. Where the special pleading comes in concerns temporal and modal embedding. The problem for SSlists is this: temporal and modal operators shift the circumstances of evaluation in such a way that, if SSI is true, we should expect that (in cases of temporal and modal embeddings of 'know') knowledge attributions will track whether the subject can rule out alternatives relevant in the subject's practical environment in the (temporally or modally shifted) circumstance of evaluation. But this prediction doesn't seem to pan out, as speakers are inclined to regard the same alternatives as relevant when evaluating non-embedded *and* embedded uses of 'know'<sup>33</sup>.

MacFarlane claims that one will not be inclined to say either of the following, predicted by SSI:

*Temporal embedding*: I know that I had two dollars in my pocket after breakfast, but I didn't know it this morning, when the possibility of counterfeits was relevant to my practical deliberations – even though I believed it then on the same grounds that I do now.

*Modal embedding*: I know that I have two dollars in my pocket, but if the possibility of counterfeiting were relevant to my practical situation, I would not know this – even if I believed it on the same grounds as now.<sup>34</sup>

The moral of the story, then, is that while each of the three leading competitor views does better than others in some respects, *none* of these

views can make sense of our willingness to attribute knowledge without some sort of Achilles heel.<sup>35</sup> And that's more or less MacFarlane's defence of (1) in the master argument.

### 7.2.2 Relativism to the rescue?

MacFarlane suggests that what we want is a semantics for knowledge attributions that satisfies the following desiderata – desiderata such that none of the three leading contender views can satisfy all of them:

*Alternative-variation*: It would explain how the alternatives one must rule out to count as knowing *vary with context* (otherwise, the view faces the dilemma facing insensitive invariantism, with respect to MacFarlane's conundrum).

*Alternative variation context ≠ use*: the alternatives one must rule out to count as knowing must not vary with context of *use* (otherwise: disagreement cannot be preserved, *a la* contextualism).

*Alternative variation context ≠ circumstances of subject*: the alternatives one must rule out to count as knowing must not vary with circumstances of the subject to whom knowledge is ascribed (otherwise: temporal and modal embeddings cannot be made sense of, *a la* SSI).

Here's where the relativist comes to the rescue. The first step is to preserve *alternative variation*, but by taking the *relevant alternatives to be determined by the context of assessment*.

The resulting view would agree with contextualism in its predictions about when speakers can attribute knowledge, since when one is considering whether to make a claim, one is assessing it from one's current context of use. So it would explain the variability data as ably as contextualism does, and offer the same way of rescuing Closure from the challenge posed by the conundrum. But it would differ from contextualism in its predictions about truth assessments of knowledge claims made by other speakers, and about when knowledge claims made earlier must be retracted. Moreover... it would vindicate our judgments about disagreement between knowledge claims across contexts. (MacFarlane 2014, 188)

What about the temporal and modal embedding problem that faced SSI<sup>36</sup>? Relativism dodges this because a parameter for a set of contextually relevant alternatives is added to the index as a parameter distinct

from world and time indices such that shifting the world and time indices (e.g. as when 'knows' is temporally or modally embedded) does *not* involve shifting also the relevant alternatives parameter (*ibid.*, 188).

Here is MacFarlane's relativist postsemantics<sup>37</sup> for "knows":

**Relativist postsemantics for "knows":** A sentence *S* is true as used at context  $c_1$  and assessed from a context  $c_2$  iff for all assignments  $a$ ,

$$\llbracket S \rrbracket_{w_{c_1}, t_{c_1}, S_{c_1}, a}^{c_1} = True$$

where  $w_{c_1}$  is the world of  $c_1$ ,  $t_{c_1}$  is the time of  $c_1$ , and  $S_{c_1}$  is the set of possibilities relevant at  $c_2$ . On this view, as he puts it:

the relation "knows" expresses does not vary with the context – there is just a single knowing relation – but the extension of that relation varies across relevant alternatives. As a result, it makes sense to ask about the extension of "knows" only relative to both a context of use (which fixes the world and time) and a context of assessment (which fixes the relevant alternatives). (*Ibid.*, 189)

MacFarlane takes the view he's proposed as one that escapes the sceptical conundrum while threading the gauntlet so as to avoid the disagreement problem that faces contextualists and the temporal and modal embedding problem that faces SSI. We can see now why MacFarlane thinks his view has all the advantages and none of the disadvantages.

## 7.3 Objections

### 7.3.1 First objection: primary relevance, environmental luck and relevant alternatives

Let's suppose Alan can tell, just by looking, a chaffinch from a goldfinch by the colours distinctive of each. Alan, however, cannot tell, just by looking, a chaffinch from a *hologram* chaffinch. Consider now these three cases:

*Case 1:* Alan is in a friendly environment (no holograms around) and sees a chaffinch and forms the belief 'There is a chaffinch'.

*Case 2:* Alan is in a friendly environment (no holograms around) and forms the belief 'There is a chaffinch'. Subsequently, Adrian tells Alan a lie: that there are hologram chaffinches mixed with the real chaffinches.

*Case 3:* Alan is in an environment where there are hologram chaffinches mixed with the real chaffinches, but he thinks he is in a friendly environment. He sees what looks like a chaffinch and forms the belief 'There is a chaffinch'.

In Case 1 all is well, and plausibly, Alan knows that what he sees is a chaffinch, so long as he can *discriminate* between chaffinches and other 'non-chaffinch' items he would *normally* encounter (e.g. goldfinches, other birds, trees, etc.). In Case 1, it doesn't matter that Alan can't discriminate between chaffinches and hologram chaffinches.

Case 2 is different in this respect. Even though there *aren't* in fact any hologram chaffinches in the local environment, there is a related epistemically 'bad' thing: a *liar* (Adrian) who *says* that there are hologram chaffinches in the area. Plausibly, Alan must now be able to rationally dismiss the hologram chaffinch alternative if he is to retain knowledge of the target proposition. Of course, if Alan were an expert, so skilled that he can spot the minute, tell-tale signs of a hologram chaffinch (this would also involve exceptional eyesight), then Alan would be able to rationally dismiss the hologram chaffinch alternative raised by Adrian *via discriminatory* epistemic support, support of the same kind by which he knows in Case 1 – *viz.*, by being able to discriminate between chaffinches and goldfinches on the basis of their appearance. Suppose he is not such an expert with remarkable eyesight, though. Even so, just because Alan can't discriminate between chaffinches and hologram chaffinches simply by the way they look, it *doesn't follow* that he couldn't (at least, potentially) *rationally dismiss* the hologram chaffinch alternative in Case 2 and thus be able to retain his knowledge that what he's looking at is a chaffinch. Alan might after all have background evidence which entitles him to rationally dismiss the hologram chaffinch alternative. For example, Alan might have background evidence that (i) the only machine which creates hologram chaffinches is in the shop being repaired; and that (ii) someone who meets Adrian's visual description is known to be prowling about the area misleading people. With this kind of background evidence, Alan can plausibly rationally dismiss the hologram chaffinch alternative *despite* lacking discriminatory support – *viz.*, despite lacking the ability to tell the difference between a chaffinch and a hologram chaffinch, on *just the basis of how they look*. Following Pritchard (2010) and Carter & Pritchard (2015), call this other kind of epistemic support *favouring epistemic support*.

Following the terminology introduced in Pritchard (2010) and further developed in Carter & Pritchard (2015),<sup>38</sup> let's distinguish (with reference

to Case 1 and Case 2) between *two distinct ways* an alternative can become epistemically relevant, such that, in order to retain one's knowledge of the target proposition, one must be able to (e.g. either by discriminatory or favouring epistemic support) rationally dismiss the alternative.

Call an alternative *primary relevant* if it is the kind of alternative that might plausibly occur in one's environment, and call an alternative *secondary relevant* if it is made relevant in some other way such as by one's becoming aware of the alternative. In Case 1, the goldfinch alternative is *primary relevant*, and in Case 2, the hologram chaffinch alternative is *secondary relevant*.

What, then, about *Case 3*? Here we have a *primary relevant* alternative; the hologram chaffinch alternative is relevant not because there's someone raising the possibility, making Alan aware of it, but for the simple reason there actually *are* hologram chaffinches mixed about with the genuine ones. Although Alan is in fact looking directly at a real chaffinch, his local environment is such that there are very close near-by possible worlds in which what he is looking at just now (and believing to be a chaffinch) is not a chaffinch but a hologram chaffinch. Alan in this case (Case 3) *can't* rationally dismiss the hologram alternative, which is primary relevant, and so fails to know.

This verdict – *viz.*, that Alan lacks knowledge in Case 3 – connects neatly with two related points. Firstly, it is to be expected that some environments are *harder* to acquire knowledge in than others. Alan is, in Case 3, in an epistemically inhospitable environment. We shouldn't expect that he should count as knowing he's looking at a chaffinch in this environment simply because he knows he's looking at a chaffinch in Case 1, where the environment is friendly (e.g., where there's no funny business going on).<sup>39</sup> Secondly, and more importantly, that Case 3 features a primary relevant alternative – one which Alan can't rule out and thus fails to know he's looking at a chaffinch – aligns with mainstream thinking about the incompatibility of *environmental* epistemic luck with knowledge.

In standard Gettier-style cases<sup>40</sup>, one plausible reason knowledge is not regarded as present is that, given the conditions of the formation of one's belief, one could very easily have been incorrect<sup>41</sup> – though in Gettier's original cases, this is on account of a 'disconnect' between the justification the agent has for the target belief, and that belief's being correct, a severed connection which is then regained only through luck.

*Environmental* luck cases are very different in this respect. Nothing *actually* goes awry. For example, in Case 3, it's not as though Alan

has good reason to believe that, say, there's a chaffinch on the tree, though in actuality what he sees is a goldfinch disguised as chaffinch, which happens to be obscuring from sight a genuine chaffinch.<sup>42</sup> Rather, the bird actually *is* a chaffinch. It's just that he could so very easily have been wrong, given the kind of *environment* he's in, one in which hologram chaffinches are in fact mixed in with the genuine article.

Putting this all together: Alan's inability to rationally dismiss the primary relevant alternative in Case 3 aligns with mainstream thinking about the relationship between knowledge and environmental epistemic luck. A kind of *ex ante* constraint on an account of knowledge is that it is going to account for why Alan *fails* to know. There are of course, some philosophers who simply insist that knowledge *is* present even in barn façade cases.<sup>43</sup> Though this is not the standard view, and at any rate, requires a wider concession, a revisionary line on the anti-luck platitude.

Against this background, I want to now sharpen what is I think at least one 'Achilles' heel' that MacFarlane's view faces – namely, that his view has trouble making sense of *primary relevant alternatives*, and that in light of this, his view generates the wrong result in cases where, intuitively, environmental epistemic luck undermines knowledge.

Consider the following case, which is a twist on Case 3:

**Case 3\*:** Alan is in an environment where there are hologram chaffinches mixed with the real chaffinches, but he thinks he his in a friendly environment. He sees what looks like a chaffinch and forms the belief 'There is a chaffinch.' Charles and Liz are in a friendly environment (no holograms around). Liz says 'Alan knows that what he is looking at is a chaffinch' and Charles evaluates this claim.

According to MacFarlane's proposal, Liz's claim that 'Alan knows that what he's looking at is a chaffinch' is true only relative to a context of use (which fixes the world/time) and context of assessment, which fixes what counts as the relevant alternatives.<sup>44</sup> As evaluated by Charles, the context of assessment fixing the relevant alternatives will be a *friendly environment*, one where hologram chaffinch alternatives needn't be ruled out – there are neither any hologram chaffinches present in Charles' environment nor has this possibility been raised, and so hologram chaffinches are not secondary relevant for Charles either. In Charles' friendly, normal environment, one can attain knowledge that one is looking at a chaffinch provided one can distinguish *chaffinches*

from *goldfinches*. And because Alan *can* distinguish chaffinches from goldfinches, MacFarlane's view rules that Liz's claim that 'Alan knows that what he's looking at is a chaffinch' comes out true as evaluated by Charles. *But it's not true!* After all, Alan is in an environment with hologram chaffinches mixed in with the genuine ones, and could very easily have pointed to a hologram rather than a genuine chaffinch and would have believed incorrectly. Moreover, Alan's belief is subject to environmental luck; he could easily have been incorrect, given the conditions of the formation of his belief, and this due to features of Alan's modal environment: in the epistemically inhospitable area where Alan is forming beliefs about chaffinches, there are very close near-by worlds in which Alan looks at a hologram chaffinch while believing he is looking at a chaffinch.<sup>45</sup>

To be clear: that MacFarlane's view fails to capture the epistemic significance of primary relevant alternatives (a point I'll revisit) is a point that does *not* depend on defending the separate though related claim that MacFarlane's view fails to capture the epistemic significance of *secondary* relevant alternatives. In fact, I think MacFarlane's view does much better in making sense of (at least one common strand of)<sup>46</sup> secondary relevant alternatives than primary relevant alternatives. To see this point, let's run a variation on Case 2 which parallels the variation we've just run on Case 3:

**Case 2\*:** Alan is in a friendly environment (no holograms around) and forms the belief 'There is a chaffinch.' Subsequently, Adrian tells Alan a lie: that there are hologram chaffinches mixed with the real chaffinches. Charles and Liz are in a friendly environment (no holograms around). Liz says 'Alan knows that what he is looking at is a chaffinch' and Charles evaluates this claim.

In Case 2\*, according to MacFarlane's proposal, Liz's claim that 'Alan knows that what he's looking at is a chaffinch' is – to reiterate – true only relative to a context of use (which fixes the world/time) and context of assessment, which fixes what counts as the relevant alternatives. As evaluated by Charles, the context relevant to fixing which alternatives count as relevant is *Charles's*. Now let's imagine two versions of Case 2\*. In the first version, suppose that everything is just as described in Case 2\*. That is, suppose that (i) Adrian is lying about the presence of hologram chaffinches to Alan, in Alan's environment, which contains no hologram chaffinches, and that (ii) Charles's environment (also, Liz's) is perfectly friendly, no holograms, and furthermore, no liars are raising

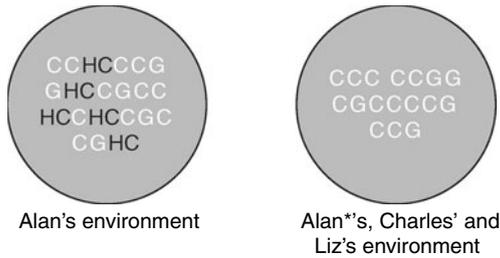
this error possibility. Now, let's run a variation on this case – call the variation Case 2\*(!). In Case 2\*(!) suppose we hold everything from Case 2\* fixed, but that we insert into *Charles's* environment a liar, (Adrian's brother, Adrian\*), where Adrian\* tells Charles that there are hologram chaffinches about. Now, when Liz says 'Alan knows that what he's looking at is a chaffinch', it's plausible that, as evaluated by Charles, our intuitions are going to shift from Case 2\* to Case 2\*(!) in at least the following respect: it's plausible to think that we are going to be *less inclined* to regard Liz's claim, as evaluated by Charles, as true when evaluated by Charles in 2\*(!) than when evaluated by Charles in Case 2\*. And this is what MacFarlane's view predicts.<sup>47</sup>

Though even if we were to grant that MacFarlane's view accommodates secondary relevance in the sense just described,<sup>48</sup> the fact that the view does not do well with primary relevance remains a pressing one. I want to now suggest that MacFarlane's inability to capture the epistemic significance of primary relevant alternatives is, in fact, an *intractable* problem for the proposal, in the same sense that MacFarlane regards temporal and modal embedding cases to be an intractable problem for subject sensitive invariantism. That is, there is no solution to be found *within* the framework being offered. In MacFarlane's case, I want to now show why the primary relevance problem is not merely a problem that MacFarlane could tweak his view in order to accommodate. It is, rather, a problem that cannot be dealt with on *any* framework on which the context of assessment is what fixes the relevant alternatives.

To emphasise this point, suppose we run yet another variation on Case 3. Call this variation Case 3\*\*.

**Case 3\*\*:** Alan is in an environment where there are hologram chaffinches mixed with the real chaffinches, but he thinks he is in a friendly environment. Alan\* is like Alan in all respects except that, in Alan\*'s environment, there are no hologram chaffinches mixed in with the real chaffinches. In their respective environments, Alan and Alan\* see what looks like a chaffinch and each forms the belief 'There is a chaffinch.' Charles and Liz are, like Alan\*, in a friendly environment (no holograms around). Liz says 'Alan and Alan\* know that what they are looking at is a chaffinch' and Charles evaluates this claim.

To represent Case 3\*\* more clearly, let 'C' represent 'chaffinch', 'G', goldfinch, and 'HC' 'hologram chaffinch.'



Now, when Liz claims that ‘Alan and Alan\* know that what they are looking at is a chaffinch’ consider what, on MacFarlane’s proposal, is going to make this true as evaluated by Charles. For one thing, on MacFarlane’s framework, the set of alternatives that are *relevant* in evaluating the truth of Liz’s assertion, as evaluated by Charles, are the set of alternatives relevant to *Charles’s* situation, which is a friendly environment, one where one must be such that one can rationally dismiss the goldfinch alternative, and not the hologram chaffinch alternative. With this in mind, consider that Liz’s assertion can be understood as the claim that Alan and Alan\* both have the following property  $[\lambda x. x$  knows that  $x$  is looking at a chaffinch].<sup>49</sup> On MacFarlane’s proposal, it comes out true, as evaluated by Charles, that Alan\* has this property. Alan\* can discriminate between chaffinches and goldfinches. So one conjunct of Liz’s claim, the part about Alan\*, comes out true. What about the other one, the part about Alan? Well, notice that the matter of whether Alan has the property  $[\lambda x. x$  knows that  $x$  is looking at a chaffinch] is, as evaluated by Charles, such that the *very same alternatives are relevant*, the alternatives fixed by the context of (Charles’s) assessment. It was already stipulated that Alan and Alan\* do not differ in their discriminatory abilities: each can distinguish chaffinches from goldfinches but not chaffinches from hologram chaffinches. But since Alan can discriminate chaffinches from goldfinches, then, Liz’s attribution to Alan the property  $[\lambda x. x$  knows that  $x$  is looking at a chaffinch] is true as evaluated by Charles *no less than* is her attribution to Alan\* the property  $[\lambda x. x$  knows that  $x$  is looking at a chaffinch]. Thus, on MacFarlane’s view, Liz’s assertion that Alan and Alan\* know they are looking at a chaffinch is true, as evaluated by Charles.

Now, again, this is the *wrong* result. And the underlying problem is the *kind* of framework MacFarlane’s offered. If the context of assessment fixes

the relevant alternatives, then in Case 3\*\* there is no basis on which we could explain how Alan\* knows but not Alan. And this is no good, given that Alan\* plausibly does know he's looking at a chaffinch, and Alan does not. (After all, Alan is in a bad environment.) Put more abstractly: the problem is that while what fixes the secondary relevant alternatives (at least, the kind explored in Case 2\*) could at least in principle be fixed by the context of assessment, primary relevant alternatives are different: what *makes* the hologram chaffinch alternative relevant for Alan in Case 3 (and its variants) is that his local environment is such that there *are in fact very close near-by possible worlds in which what he is looking at just now (and believing to be a chaffinch) is not a chaffinch but a hologram chaffinch*. Put another way: what makes the hologram chaffinch alternative relevant for Alan in Case 3 (and variants) is Alan's *modal environment, as determined by Alan's local environment*. And Alan's *modal environment* remains the same across all possible contexts of assessment.<sup>50</sup> This is structural problem. If contexts of assessment are to be understood as what fixes the relevant alternatives, then primary relevant alternatives can be ignored, and ignored in a way that (in cases like Case 3\* and 3\*\*) reveals how MacFarlane's view will count environmentally lucky beliefs as knowledge.

### 7.3.2 Second objection: secondary relevance and normatively relevant alternatives

A clarificatory point will be helpful in transitioning from the argument outlined in the previous section to the one I'll be advancing in this section. The first point is to make clear is that the argument just outlined does not *ultimately* depend for its efficacy against MacFarlane's view, on any antecedent acceptance of the thesis that environmental luck is incompatible with knowledge<sup>51</sup> (even though I think this insight is entailed by the fundamental anti-luck insight, which should guide our theorising about knowledge).<sup>52</sup> To emphasise, the objection highlighted in the previous section can be framed in a way that *avoids talk of epistemic luck entirely*. The 'luck-free' formulation of the objection is that (with reference to case 3\*\*) the same alternatives are relevant (keyed to Charles' context of assessing Liz's claim about Alan and Alan\*) in evaluating both Alan's and Alan\*'s claims to know there is a chaffinch. *But*, as the argument goes, there are primary relevant alternatives Alan must rule out to know (given the hologram chaffinches in Alan's environment) which Alan\* needn't rule out. MacFarlane's view thus (and again, without using the language of luck) generates the view that, as assessed

by Charles, Alan knows if Alan\* knows. And this is bad because, unlike Alan\*, Alan needs to be able to dismiss the hologram chaffinch alternative to plausibly count as knowing, something he can't do.

To formulate this point more cleanly, and abstractly, just take a case,  $C_1$ , where a set of alternatives  $A$  are primary relevant and another case  $C_2$  where another *distinct* set,  $A^*$ , are primary relevant. Our view of knowledge should say that whether there's knowledge in  $C_1$  depends on whether the subject in  $C_1$  can rule out  $A$  and whether there's knowledge in  $C_2$  depends on whether the subject in  $C_2$  can rule out  $A^*$ . Now imagine a context of assessment where, by whatever mechanisms MacFarlane posits, a set  $A^{**}$  – distinct from  $A$  and from  $A^*$  – are relevant. Because MacFarlane claims that the truth of the assessment (*vis-à-vis*  $C_1$  and  $C_2$ ) depends on whether the subjects can rule out  $A^{**}$ , it follows that there's knowledge in  $C_1$  if, and only if, there's knowledge in  $C_2$ . But any good theory of knowledge tells us this biconditional is false.<sup>53</sup> Thus, and without reference to epistemic luck, MacFarlane's view has a structural problem with primary relevant alternatives, and thus with cases like  $3^*$  and  $3^{**}$ .

I want to now outline a second style of objection, one which turns neither on environmental epistemic luck nor primary relevant alternatives. Recall that *secondary relevant* alternatives are alternatives which are not made epistemically relevant in virtue of their *being* the sort of things one finds in the environment one is in. Recall again, for reference, Case 2. In Case 2, even though Alan was in a friendly environment (one with no holograms chaffinches around, only chaffinches and goldfinches, which Alan can tell apart), nonetheless the hologram chaffinch alternative became epistemically relevant given that Adrian *said* to Alan that there are hologram chaffinches mixed with the real chaffinches (and, even though this wasn't true).

In Case 2, the hologram chaffinch alternative becomes secondary relevant once Alan is in fact aware of the alternative. In short, the idea is that once a subject becomes aware of an alternative *qua* alternative,<sup>54</sup> and recognises its incompatibility with what one believes, one must be able to rationally dismiss this alternative (e.g. by either discriminatory or favouring epistemic support) if she is to retaining her knowledge of the target proposition.

In recent co-authored work,<sup>55</sup> I've argued that there are *two* very different ways that an alternative can become *secondary* relevant. And as it turns out, it's the *second kind of secondary relevance* that's going to be problematic for MacFarlane. Compare now Case 2 (featuring the 'first'

kind of secondary relevance, where an alternative becomes relevant by one's becoming *aware* of the alternative qua alternative) with Case 4:

**Case 4:** The zoo that Zula is visiting has a number of signs posted near the zebra enclosure which state (falsely) that the creatures therein are not zebras but cleverly disguised mules. Suppose further that Zula *should* have spotted these signs, but fails to simply because she is a very inattentive person. Had Zula been attentive, then she would have been made aware of the cleverly disguised mule alternative, and it would have been incumbent upon her, if she is to retain her knowledge that the creature before her is a zebra, to adduce favouring epistemic support which excludes this alternative.<sup>56</sup>

*Question:* Does the mere fact that Zula has failed to spot this (misleading) counterevidence suffice to ensure that this epistemic demand – *viz.*, that she must be able to rationally dismiss the cleverly disguised mule alternative in order to retain knowledge of the target proposition – is not placed upon her? In short, no.<sup>57</sup>

To appreciate this point, consider here the precedent we find in the literature on epistemic defeat – namely, that both *psychological* and *normative defeaters* are capable of defeating one's epistemic justification.<sup>58</sup> The simple picture, following Lackey (2010, 317) goes, like this: *S*'s (would-be knowledge or justification) that *p* can be defeated either *psychologically* or *normatively*, in virtue of *S* in the former case *possessing* some belief, *q* or, in the case of normative defeat, being such that *S* (epistemically) *should believe q*, where *q* is a proposition the truth of which would in either rebut (count against the truth of) or undercut (count against the reliability of the formation of) *p*. Defeaters themselves can be defeated by other defeaters. However, when they are not, (undefeated) defeaters vanquish knowledge and justification. This is the simple story, and it's one on which facts about what I epistemically should believe are epistemically significant with respect to what I count as knowing.

If we were to say the *only* kind of secondary relevant alternatives which mattered (i.e. in the sense that: they are epistemically significant such that rationally dismissing them is required for retaining knowledge of the target proposition) were ones that are relevant in virtue of awareness, then we would be in a very awkward spot: one where embracing the orthodox distinction between psychological and normative defeaters would be unprincipled.

Of course, one might try to preserve consistency by (along with insisting that the only kinds of secondary relevant alternatives are ones of which

one is in fact aware) embracing to boot a mad dog ‘rejectionist’ view of normative defeat. On such a view, one claims that the *only* epistemic requirements that hold of a given subject are requirements to believe in accord with the evidence, and further, that in every case the relevant evidence includes all and only the evidence that the subject has in her possession at the time in question. Following Goldberg (2015), we can call this view *Extreme Cliffordism*.<sup>59</sup> If one has reason to think Extreme Cliffordism is true, then one might have cause to be a rejectionist about normative defeat (and as such, would be in a position to take a *principled* line against the view that, in Case 4, the cleverly disguised mule is a relevant alternative for Zula because she epistemically should have had evidence (albeit, misleading) which she overlooked).

But Extreme Cliffordism is in fact very hard to swallow. Consider for instance, what Goldberg (2015) has to say about the view:

We might think e.g. of the case of the physicist who does not attend to the criticisms of his peers, and so who continues to believe that his theory is well-supported when in fact there is a good deal of recently-discovered counterevidence (Kornblith 1983, p. 36); the doctor who failed to keep up with the medical journals she ought to be keeping up with (H. Smith 1983, pp. 544–545); the daycare provider who fails to read the allergy report of the new child in his care (Smith 1983, p. 551); the person who fails to appreciate that the situation she confronts calls for a special effort on her part (Montmarquet 1992, p. 339); the situation in which a person who unjustly benefits from the *status quo* rejects the case for the injustice of things out of “affective ignorance” (Moody-Adams 1994, p. 296); or the man who fails to glimpse at the note on the fridge prior to assuming that there is enough milk for breakfast (Gibbons 2006, p. 22). Let us stipulate that in each case, the subject’s belief conforms to the evidence (s)he currently has in her/his possession. In that case, the proponent of [sic. Extreme Cliffordism] will hold that *the subject’s belief is justified in all of these cases*. (Goldberg 2015, §5, my italics)

There are further problems with Extreme Cliffordism, beyond just that it generates the consequences Goldberg highlights, but I won’t review them here.<sup>60</sup> It suffices for the point I want to make here that there is, plausibly, a variety of secondary relevant alternative which is such that MacFarlane has trouble making sense of it. And we needn’t positively establish the falsity of Extreme Cliffordism to accept that there is already a kind of philosophical precedent for thinking in a way that

is compatible with a negative answer to the question posed following Case 4 about Zula. Accordingly, I think it's very plausible that what matters for secondary relevance is *either*

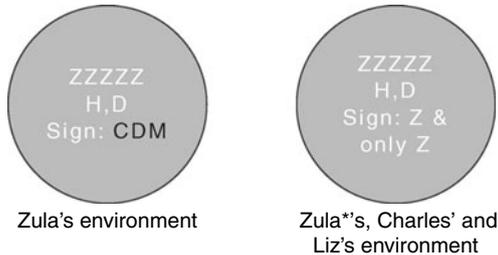
- (i) that the subject is aware of the alternative qua alternative (e.g. Case 2)
- (ii) *or* that this is an alternative that she (epistemically) ought to be aware of (e.g. Case 4).

Either condition suffices to place the additional epistemic demand on the subject.<sup>61</sup> Call (i)-style secondary relevance *awareness secondary relevance* and (ii)-style secondary relevance *normative secondary relevance*. The point of Case 2\* was to concede to MacFarlane that there is at least no obvious barrier to thinking of *awareness secondary relevant* alternatives as being fixed by the context of assessment.<sup>62</sup> However, let's consider why things are much more problematic for MacFarlane in the case of *normative secondary relevance*. Consider the following twist on Case 4:

**Case 4\*:** The zoo that Zula is visiting has a number of signs posted near the zebra enclosure which state (falsely) that the creatures therein are not zebras but cleverly disguised mules. Suppose further that Zula *should* have spotted these signs, but fails to simply because she is a very inattentive person. Zula\*, let's suppose, is like Zula in all respects (e.g. she has the same discriminatory abilities – she can tell zebras from things that might plausibly be found in a zoo, e.g. horses, moose, but *not* from cleverly disguised mules) except that she's in an environment where there is *no* misleading sign, but rather (to make things simple) an accurate sign which states that there are 'zebras and only zebras' in the zebra enclosure. Suppose further that Zula\* is attentive and does see this sign. Now suppose that Zula and Zula\* both look at a (genuine) zebra in their respective environments and form the belief 'There is a zebra'. Charles and Liz are, like Zula\*, in a friendly environment (no misleading signs around). Liz says 'Zula and Zula\* know that what they are looking at is a zebra' and Charles evaluates this claim.

In the pictorial representation below, let 'Z' represent a zebra, and let 'CDM' represent the misleading message 'cleverly disguised mules in the area', which is on the sign by the zebra enclosure in Zula's environment, which she inattentively overlooks. And let 'Z&OnlyZ' represent the non-misleading message that there are *Zebras and only zebras in the enclosure*

stated on the relevant sign in Zula\*'s, Liz's and Charles environment. To make things simple, suppose that the only other animals in either zoo are horses 'H' and deer 'D'.



Now, when Liz claims that 'Zula and Zula\* know that what they are looking at is a zebra' let's again consider what, on MacFarlane's proposal, is going to make this true, as evaluated by Charles. For one thing, on MacFarlane's framework, the set of alternatives that are *relevant* in evaluating the truth of Liz's assertion, as evaluated by Charles, are the set of alternatives relevant to *Charles's* situation, which is a friendly environment, one where one must be such that one must be able to distinguish zebras from horses and deer, but not from cleverly disguised mules.

In a way analogous to what we saw in Case 3\*\*, Liz's assertion (in Case 4\*) can be understood as the claim that Zula and Zula\* both have the following property [ $\lambda x. x$  knows that  $x$  is looking at a zebra]. On MacFarlane's proposal, it comes out true, as evaluated by Charles, that Zula\* has this property. Zula\* after all can discriminate between zebras and horses and deer, the kinds of alternatives one must be able to rationally dismiss in Zula\*'s and Charles's environment in order to know one is looking at a zebra. And moreover, there is nothing Zula\* has overlooked out of laziness that would have suggested anything otherwise. So one conjunct of Liz's claim, the part about Zula\*, comes out true, as evaluated by Charles. What about the other one, the part about Zula?

Well, notice that the matter of whether Zula has the property [ $\lambda x. x$  knows that  $x$  is looking at a zebra] is, as evaluated by Charles, such that the *very same alternatives are relevant*, the alternatives fixed by the context of assessment, Charles & Zula\*'s environment. It was already stipulated that Zula and Zula\* do not differ in their discriminatory abilities. Each can distinguish zebras from goldfinches and horses and deer but not from cleverly disguised mules. But since Zula\* *can* discriminate zebras from horses and deer, then, Liz's attribution to Zula the property

[ $\lambda x. x$  knows that  $x$  is looking at a zebra] is true as evaluated by Charles *no less than* is her attribution to Zula\* the property [ $\lambda x. x$  knows that  $x$  is looking at a zebra]. Thus, on MacFarlane's view, Liz's assertion that Zula and Zula\* know they are looking at a zebra is true, as evaluated by Charles.

But problematically for MacFarlane, there is a *normatively secondary relevant* alternative for Zula which is glossed over when Charles' environment is what's regarded as fixing which alternatives are relevant. But since it is a *normative* requirement that underwrites the epistemic significance of cleverly disguised mule alternative for Zula, how is this to be preserved, as evaluated by Charles, since *Charles has no such normative* requirement: there is no such misleading sign in Charles' environment such that he should have seen it!

#### 7.4 Concluding remarks

MacFarlane's master argument for a truth-relativist semantics for knowledge attributions is that it preserves the best advantages of other leading proposals: that is, it can (unlike the insensitive invariantist) explain variability data without denying closure), and it can do this without losing disagreement, as the contextualist does, or getting the wrong results in temporal and modal embedding cases, as the SList does. But I hope to have shown that MacFarlane's view has its own Achilles' heel. Many of the advantages MacFarlane's view is able to achieve over the other competing accounts of the semantics of knowledge attributions are advantages gained, specifically, by keying relevant alternatives to a context of assessment which can vary while the context of use and associated state of the world remain fixed.

But this very move, as I've argued, will leave MacFarlane in a position where he, like the alternatives he's criticized, must do his own special pleading. By keying relevant alternatives to the context of assessment, MacFarlane's view is going to get the wrong result in cases featuring *primary relevant alternatives* (e.g. 3\*\*) and cases featuring *normative secondary relevant alternatives* (e.g. 4\*). An associated cost with getting the wrong result in cases like 3\*\* is that MacFarlane fails to preserve the insight that environmental epistemic luck is incompatible with knowledge; though more fundamentally the problem is that the view fails to make sense of the epistemic significance of primary relevant alternatives. An associated cost with getting the wrong result in cases like 4\* is that the MacFarlane's view stands in tension with ordinary thinking about normative defeat.

I want to conclude by considering, and responding to, one way that a proponent of assessment-sensitive knowledge attributions might attempt to evade the Achilles' heel I've outlined. In MacFarlane's first presentation of the argument (2005) for an assessment-sensitive treatment of 'knows', he says that what is determined by the context of assessment is not which alternatives are relevant, *per se*, but more generically, the *epistemic standard*. A relativist retreat seems available, along the following lines: (i) the arguments adduced in this chapter were problems associated with keying relevant alternatives to the context of assessment; (ii) MacFarlane's more generic 2005 view did not do this, specifically; what was keyed to the context of assessment was simply the 'epistemic standard' (iii) therefore, MacFarlane's 2005 remains unscathed from the objections.

The problem with this argument is that the *source* of the problems, problems framed in terms of relevant alternatives, posed to a view articulated in terms of relevant alternatives, was a rotten core – no matter how we fill in the substantive details of an epistemic standard (e.g. whether satisfying a standard required for knowledge involves ruling out relevant alternatives, satisfying an ability condition, a safety condition, etc.) problem cases materialise. To see why, suppose for a moment that we *quit talking about relevant alternatives altogether* and opt for a different way of thinking about what satisfying an epistemic standard requisite for knowledge involves on the part of the knower.

Here's one very elegant such view, one defended by Linda Zagzebski (1996), John Greco (2010, 2012) and Ernest Sosa (2007; 2009; 2011; 2015). Because this view is (like a relevant alternatives view) simple and elegant, it will be helpful for illustrative purposes (and the point I'm making does not turn at all on the view's being correct). The view I have in mind here is *robust virtue epistemology* (RVE). In short:

**Robust virtue epistemology (RVE):** S knows that *p* iff S's truly believing *p* is because of the exercise of S's intellectual virtue(s).<sup>63</sup>

Suppose that RVE gives the right view about what is *involved* on the part of the knower in satisfying the epistemic standard requisite for knowing. Again, this is contentious, but let's run with it for the moment. Now, imagine two versions of a testimony case:<sup>64</sup>

**Testimony-1:** The environment is an epistemically inhospitable one. Many people around are both dishonest and confused about where the Sears Tower is (while appearing honest and non-confused), and

many are up for playing a malevolent joke. Jenny asks the nearest passer-by, Bruno (who looks reasonably credible, not obviously drunk) directions to the Sears tower, and he, in fact the only reliable and benevolent testifier in the area, provides her with accurate directions.

**Testimony-2:** The environment is an epistemically hospitable one. No one around is lying or confused about where the Sears Tower is, and no one is up for playing a malevolent joke. Stipulate that everyone around looks on appearance as in Testimony-1. Jenny\* asks the nearest passer-by, Bruno\* (like Bruno in all respects), directions to the Sears tower, and he provides her with accurate directions.

On the assumption that we should fill out what's involved in satisfying an epistemic standard on the part of the knower is *as the RVE proponents tells us*, watch how MacFarlane's proposed semantics quickly runs in to a jam. Just suppose Liz and Charles are in the friendly environment of *Testimony-2*: Liz says 'Jenny and Jenny\* know where the Sears Tower is at 't', a time after each receives her respective testimony. As assessed by Charles, in the friendly environment, MacFarlane's 2005 view tells us that the standard that must be satisfied for Liz's claim to be true is *the standard that must be satisfied in Charles', Liz's and Jenny\*'s environment*. In this environment, if we fill out the epistemic standard as RVE does, the standard is met so long as one's believing truly is *because of one's intellectual virtues (V)*, where (in Testimony 2) *exercising V in a way that suffices for knowing* just amounts to being careful enough to not ask someone who is visibly drunk, who looks deceptive or shady qua testifier. But if *that's* what's involved in satisfying the epistemic standard required for knowing, then, as evaluated by Charles, both conjuncts of Liz's attribution to Jenny and Jenny\* of having the property [ $\lambda x. x$  knows that  $x$  where the Sears tower] come out true. But obviously Jenny doesn't know in Testimony-1.

Although I framed the previous example, on the assumption that RVE is true (in order to make a limited point – that on some notable views in mainstream epistemology, relativising an epistemic standard that must be satisfied for knowledge to the context of assessment generates bad results), I want to conclude by noting a more general point, one which is evidenced by each kind of example I've suggested to be problematic for an assessment-sensitive semantics for knowledge attributions.

MacFarlane's approach of relativising the standard that must be satisfied for knowing to a context of assessment (either in the specific way

of 2014 or the generic way of 2005) is incompatible with a general view in epistemological meta-theory about the role that extra organismic elements of one's environment can play in the acquisition of knowledge. This view is called *epistemic anti-individualism*: the view that (in short) what converts true belief to knowledge can supervene at least partly on elements of one's local and/or modal environment.<sup>65</sup>

A rationale for why MacFarlane's view gets the wrong results, in Case 3\*\* and 4\* and in the testimony case just considered, can be neatly recast in terms of epistemic anti-individualism: by making the *context of assessment* the relevant context, one abstracts away from the environment of the subject of the knowledge attribution in a way that rules out epistemic anti-individualism in any case where the environment of the subject of the knowledge attribution and the context in which the knowledge attribution is assessed for truth/falsity *must* be kept apart.<sup>66</sup>

I want to conclude this chapter by raising a terse dilemma, one which connects the first part of the chapter with the second. If MacFarlane is right that assessment sensitivity is the *mark of relativism*, it's hard – at least, in light of industry standard epistemological criteria – to see how *knowledge* is relative.

# 8

## New Relativism: Epistemic Aftermath

*Abstract.* This chapter begins by qualifying the epistemologically oriented objections raised in the previous chapter against MacFarlane's (2005; 2009; 2014) 'new epistemic relativism'. These objections are conceded not to be ultimately *decisive* against the view unless we can further defend them to be *worse* theoretical costs than the costs that are incurred by other competing semantics for knowledge attributions. Rather than to attempt to establish this point about comparative costs, I opt in this chapter for an entirely different kind of rationale on the basis of which it will be argued that any import new epistemic relativism might lay claim to having in epistemology is ultimately undermined. The rationale advanced to this end involves putting together several pieces, though the overarching move can be stated simply: I suggest, in a fashion that draws some close parallels with Allan Hazlett's (2010) recent work on knowledge, factivity and knowledge ascriptions, that the *more* compelling MacFarlane's argument is for his conclusion that the ordinary concept 'knows' is assessment-sensitive, the more reason the *epistemologist* has for thinking that the ordinary concept of knowledge is *epistemologically* uninteresting, and moreover, so are ordinary knowledge ascriptions. Crucial to the line advanced here will be an examination of what contemporary epistemology might look like if it were centred around an assessment-sensitive concept of knowledge.

### 8.1 Semantic stalemate?

If the epistemologically driven arguments in the previous chapter against an assessment-sensitive semantics for knowledge attributions are sound, then we can add relativism to the list of views which do well in some respects in explaining our practices of attributing knowledge but which

ultimately succumb to some kind of Achilles' heel. On the Achilles' heel front, the highlights for the main competitors to MacFarlane-style relativism were that: insensitive invariantism can't very well explain the variability of our willingness to attribute knowledge.<sup>1</sup> Contextualism has trouble making sense of disagreement.<sup>2</sup> And SSI gets the wrong result in cases where 'knows' occurs under temporal and modal embeddings.<sup>3</sup> Let's grant MacFarlane that his assessment-sensitive view can lay claim to at least as many advantages as any of the other views can. The problem is, MacFarlane's view (as we saw in Chapter 7) has its own Achilles' heel in the form of three related epistemological problems:

*Primary Relevance Problem:* The view fails to capture the epistemic significance of *primary* relevant alternatives which are relevant by obtaining in the subject's local environment (e.g. Case 3\* and 3\*\*); while this is a problem in its own right, a further implication is that the view generates the heterodox result that environmental epistemic luck is compatible with knowledge;<sup>4</sup>

*Normative Secondary Relevance Problem:* The view fails to make sense of *normative secondary relevant* alternatives (e.g. Case 4\*), a problem that seems avoidable only via a wholesale rejection of normative epistemic defeaters – and thus by retreating to a kind of implausible version of 'Extreme Cliffordism'.

*Epistemic Anti-Individualism Problem:* More generally, the view was shown to rule out, *ex ante*, a plausible view in epistemological meta-theory called *epistemic anti-individualism*, a view on which what converts true belief to knowledge can supervene at least partly on an agent's local (and/or modal) environment (e.g. Kallestrup & Pritchard 2011, 2012, 2013).<sup>5</sup>

Though I take the previous chapter to have highlighted how the *primary relevance*, *secondary relevance* and *epistemic anti-individualism* problems are serious epistemological concessions to have to make for one who opts for an assessment-sensitive semantics for knowledge ascriptions, it is responsible to qualify things now. Here is a very difficult philosophical question: call it the *Achilles' Heel Question*:

*Achilles' Heel Question:* Are the three central epistemological problems facing MacFarlane's view *worse* than the Achilles' heels claimed by MacFarlane to face insensitive invariantism, contextualism and SSI?<sup>6</sup>

Trying to answer the Achilles' Heel Question in the affirmative strikes me as walking in to a trap. It's unclear after all how one is to go about weighing objections grounded in such different kinds of considerations. Obviously, the objections raised in the previous chapter are objections situated around considerations that matter in *epistemology*. Should such considerations have *special* weight, given that 'knows' is central to epistemology? Alternatively, should they have *less* weight, given that epistemologists might be interested in something under the description of 'knowledge' which differs from the way that the ordinary folk (in the sense of Kusch (2010))<sup>7</sup> use 'knows'?<sup>8</sup>

Unless one wants to attempt to dismiss MacFarlane-style relativism about knowledge attributions (hereafter, for ease of reference call this 'new epistemic relativism') by giving a positive answer to the Achilles' Heel Question (again, a task that looks like a trap), we reach an interesting kind of 'gambit point' with respect to the *plausibility* of epistemic relativism, broadly construed.

Here is the big picture: a recurring theme from Chapters 3–5 was (to put it crudely) that traditional arguments for epistemic relativism run in to – among other problems – a wall at the last hurdle: even if we grant that the Pyrrhonian, Dialogic and Incommensurability/Circularity style arguments surveyed could overcome the various objections raised which were specific to each argument-type, it remains – I argued – that none of these arguments would ultimately be successful in recommending relativism *over* scepticism.<sup>9</sup> It is natural to think, then, that new epistemic relativism might as well face the same hurdle and thus that we might feasibly be in a position to 'reverse engineer' an argument against new-epistemic relativism motivated by the track record of the more traditional versions: that is, we might attempt to show why new epistemic relativism is going to inevitably face the same kinds of problems faced by the Pyrrhonian, Dialogic and Incommensurability/Circularity-style arguments.

But the problem is that it's not at all clear why it *would*. Here a representative example will be helpful. Take for instance the manoeuvre we find in the strand of dialogic argument for epistemic relativism, which Siegel called the 'no-neutrality, therefore relativism' argument – the critical focus of Chapter 4, and which took as a starting point the dispute between Galileo and Bellarmine. While I argued that 'actualist' and 'possibilist' versions of the 'no neutrality, therefore relativism' argument were ultimately unworkable, it was shown that even if we *granted* to the proponent of such an argument that the problems raised could be overcome, the *most* that the proponent of this argument for

epistemic relativism could plausibly lay claim to was a kind of 'intermediate conclusion' to the effect that it's *not* the case that there can be a non-relative resolution of the dispute concerning the existence of the moons. And from this position, it was shown that a decisive motivation for relativism over scepticism was simply lacking. And, more generally, this turned out to be a familiar kind of gambit point for the would-be relativist.<sup>10</sup>

The proponent of assessment-sensitive knowledge attributions at this point might ask 'And how is this relevant?' And indifference here is fair enough. *Traditional* arguments for relativism have proceeded by taking a least a few steps in tandem with the sceptic before jumping ship. MacFarlane, however, doesn't take *any* steps in tandem with the sceptic *en route* to his version of epistemic relativism. In fact, the only remark MacFarlane passes on the topic of scepticism is to briefly regard it as a bad view to be stuck with, when commenting on what he took to be the perils of insensitive invariantism.<sup>11</sup> 'Reverse engineering' an argument against new relativism inspired by the track record of *traditional* forms of epistemic relativism thus doesn't look promising.

But if *this* is right, then here seems to be the nub of things: *unless* we want to attempt to defend a positive answer to the Achilles' Heel Question, the most that can responsibly be laid claim to on the basis of the epistemologically oriented objections to new epistemic relativism outlined in the previous chapter is that, *contra* MacFarlane, *it's not clear that we have positive reasons to prefer a relativist semantics for knowledge attributions over the key competitors*. That is: we can say to the new relativist that the burden remains with her to show how (in light of the epistemologically oriented objections raised) she thinks she can claim a decisive win rather than a mere stalemate against any of the other competing semantics for knowledge attribution. And, if this is right, then we see a sense in which MacFarlane's brand of new relativism constitutes a form of epistemic relativism that is considerably more difficult to dismiss than the traditional forms.

Some might find this position unsatisfying. The thought that we must settle with granting the proponent of an assessment-sensitive semantics for 'knows' a stalemate (*vis-à-vis* competing semantics for knowledge attributions) looks like a significant concession indeed in the relativist direction. This is, I think, all the more reason for mainstream epistemologists to take seriously the kind of challenge to orthodoxy (in epistemology) posed by new relativism, and this so even if mainstream epistemologists are right to by and large proceed as though more *traditional* forms of relativism (e.g. as discussed in Chapters 3–5) are unmotivated.

Against this background, I want to sketch out what I take to be the most promising line for epistemologists who regard granting a stalemate to any semantics for knowledge attributions that succumbs to the three epistemological problems raised in the previous section to be unacceptable. Here's how the chapter will proceed. In the next section, I want to trace out the *implications* of embracing a truth-relativist semantics for knowledge attributions, more widely, in epistemological theory. The question, in short, will be one MacFarlane asks himself, rhetorically, in his 2009 defence of a relativist semantics for knowledge attributions. In particular, in his section 'Questions for the Relativist', one question he asks is: 'are there other expressions for which a relativist treatment is needed? How does "know" relate to them?'<sup>12</sup> A more specific version of this question is:

*Epistemological Ramification Question (ERQ):* if 'knows' gets relativist treatment, then since knowledge relates intimately with other epistemic concepts, do any other epistemic concepts also need a relativist treatment?

As we'll see in the next section, the dominoes quickly fall. I'll suggest that the most reasonable answer we can give to the epistemological ramification question (ERQ) implies a kind of 'wholesale epistemic relativism', such that: a truth-relativist semantics for knowledge attributions would plausibly – depending to some extent on what our substantive epistemological commitments are – require us to give up absolute truths with respect to a range of other connected epistemic notions. I'll then proceed to build upon this the idea. In particular, I suggest that if what I'll be calling *wholesale* epistemic relativism is in fact a natural consequence of embracing a truth-relativist semantics for knowledge attributions, then an avenue opens up for running an argument that parallels in some key respects a line recently taken by Allan Hazlett (2010) in his provocative paper 'The Myth of Factive Verbs'. The central argument of Hazlett's paper is that the non-factive use of 'knows' that he takes to be characteristic of the *folk* concept of 'knows' is – given that epistemologists are interested in a factive epistemic standing under the description of 'knowledge' – evidence that 'traditional epistemology and ordinary language epistemology (as we might call the theory of knowledge attributions) would both be best served by going their separate ways' (Hazlett 2010, 522).

## 8.2 Epistemic aftermath

In tracing out epistemological ramifications of a relativist treatment of 'knows' in *epistemology*, it's helpful to begin with tightest conceptual connections and move outward from there. Take, as an example case, Williamson's (2000) knowledge-evidence equivalence:  $E=K$ . Suppose, for *reductio*, that  $E=K$ , and further, that the truth-conditions for  $E$  are *not* assessment sensitive, but the truth-conditions for  $K$ , are. The resulting tension would be untenable (at best), at worst, contradictory. While of course Williamson's view is controversial<sup>13</sup>, it seems that if Williamson is right that our evidence is what we know, and thus that  $S$ 's evidence includes  $E$  if, and only if,  $S$  knows  $E$ , then the truth-relativist about knowledge ascriptions had better be prepared to embrace the view that that *evidence ascriptions* are assessment-sensitive.<sup>14</sup>

Let's move now from an equivalence thesis to a reductivist thesis. We needn't look further than the most standard contemporary version of intellectualism about knowledge-how, a view discussed in passing in Chapter 4. Reductivist versions<sup>15</sup> of intellectualism insist that knowing how to do something is just a species of propositional knowledge (Stanley 2011, 207).<sup>16</sup> As Stanley puts it:

[...] you know how to ride a bicycle if and only if you know in what way you could ride a bicycle. But you know in what way you could ride a bicycle if and only if you possess some propositional knowledge, viz. knowing, of a certain way  $w$  which is a way in which you could ride a bicycle, that  $w$  is a way in which you could ride a bicycle. (*Ibid.*, 209)<sup>17</sup>

Like Williamson's  $E=K$  thesis, Stanley's reduction of knowledge-how to a kind of knowledge-that is also controversial, though very much a live and increasingly popular view in contemporary epistemology. Suppose, for *reductio*, that knowing how to do something is (*a la* Stanley) just a kind of propositional knowledge, and further, that the truth-conditions for knowing how to do something (e.g. as in the case of attributions of the form 'Hannah knows how to ride a bike') are *not* assessment sensitive, but the truth-conditions for proposition knowledge are, such that 'Hannah knows  $p$ ' is assessment-sensitive, where  $p$  is a proposition specifying of a way  $w$  which is a way in which Hannah could ride a bicycle, that  $w$  is a way in which Hannah could ride a bicycle. Again, the resulting tension would be untenable (at best), at worst, contradictory.

Consider now another kind of popular reductivist proposal: according to one prominent account of what it is to understand why something X is the case, this doesn't involve, as Peter Lipton (2004, 30) puts it 'some sort of super knowledge, but simply more knowledge: knowledge of *causes*'.<sup>18</sup> The view that understanding-why (hereafter, understanding) is just a *kind* of propositional knowledge is reductive in the same sense as is Stanley's claim that knowledge-how is a kind of propositional knowledge. As Grimm (2014), notes, explanatory understanding reductivism has enjoyed considerable endorsement.<sup>19</sup> It is, as Grimm calls it, the 'traditional view of understanding', the first endorser of which was Aristotle. If knowledge-attributions are assessment-sensitive, then just as (for the reductive intellectualist) we must be prepared to give 'knows-how' a relativist treatment, we must also be prepared to give understanding attributions a relativist treatment. After all if 'S knows that *p*' is assessment-sensitive, then, if, as proponents of the traditional view of understanding tell us, 'S knows that *p*' gets the same truth conditions as 'S understands why *p*', where *p* is the relevant kind of causal proposition, 'S understands why *p*' will be assessment sensitive.

Let's quickly take stock: one who wishes to maintain that knowledge attributions are assessment-sensitive must be prepared to also give a relativist treatment to other epistemic notions that one also endorses as either equivalent to, or reducible to a species of, propositional knowledge. This, as we've seen, will be the case with: (i) evidence, for one who embraces a Williamsonian E=K equivalence; (ii) knowledge-how, if one embraces Stanley-style reductive intellectualist account of knowledge-how; and (iii) explanatory understanding, if one signs on to the traditional account of understanding.

So these are three very quick dominoes to fall, granted, they are to differing degrees theory-laden. On, reflection, it should not be surprising that one who wants to give 'S knows that *p*' a relativist treatment can't very well hold on to the thought that other epistemic notions that they explicitly regard as either equivalent to, or reducible to something they give a relativist treatment to could themselves be given an absolutist treatment. However, I think that the ramifications of embracing 'new epistemic relativism' with respect to knowledge attributions reaches outward quite a bit further into the epistemic landscape than just to other epistemic notions that are regarded as equivalent to, or reducible to, propositional knowledge.

In particular, consider the (albeit, complicated) relationship between knowledge and *justification*. According to traditional thinking, knowledge

is analysable into more fundamental components, and, these components are justification, truth and belief (*modulo* refinements, e.g. the addition of some other component 'x'). This is one way to capture the JTB analysis of knowledge, one where knowledge is the analysandum, and justification is among the analysans. Of course, Williamson (2000) reverses this story. But let's proceed with the traditional view. One ambiguity at play in the provision of JTB analyses concerns how to answer what Ichikawa and Steup (2014) call a certain 'metaepistemological interpretive question'. For instance, in analysing knowledge *in terms of* justification, are we taking ourselves to put forward a claim with metaphysical import to the effect that '*what it is for S to know p is for some list of conditions involving S and p to obtain?*'<sup>20</sup> Or, alternatively, is the inclusion of justification as among the analysans of knowledge rather a claim about our *concepts* of knowledge and justification, to the effect that the concept of the former includes, as a component part, the concept of the latter? Unfortunately, as Ichikawa & Steup (2014) note:

In practice, many epistemologists engaging in the project of analyzing knowledge leave these metaphilosophical interpretive questions unresolved; attempted analyses, and counterexamples thereto, are often proposed without its being made explicit whether the claims are intended as metaphysical or conceptual ones.

However, there is something of a bright side. They add that:

In many cases, this lack of specificity may be legitimate, since all parties tend to agree that an analysis of knowledge ought *at least to be extensionally correct in all metaphysically possible worlds*.

In light of this, proponents of a traditional analysis of knowledge, despite not endorsing an equivalence claim (e.g.  $K=J$ ) *nor* a reductive claim (e.g.  $K$  is itself just a variety of  $J$ ), will still find themselves in the following position: if knowledge is assessment-sensitive, then plausibly, so must justification be; for if it *weren't*, then, possibly, a claim of the form  $S$  knows that  $p$  could be true as uttered at  $c_1$  and assessed from  $c_2$ , even when ' $S$  is justified in believing that  $p$ ' is false as uttered at  $c_2$  and assessed at  $c_2$ , a result that contradicts any form of the traditional analysis that takes a  $K=JTBx$  analysis (where  $x$  is any further condition) to be extensionally correct in all metaphysically possible worlds.<sup>21</sup>

But there's a further way to argue for this point, through a side door, at least, if we restrict our focus to internalist justification. As Boghossian

(1989) puts it, internalists are widely sympathetic to a kind of ‘grasping’ condition that must be in place for an agent to be justified in believing a proposition  $p$  on the basis of another proposition  $q$ , where  $q$  is the agent’s evidence for  $p$ :

*Grasp*: A knower must grasp the connection between the evidence and what it is evidence for if his belief is to be justified.<sup>22</sup>

Boghossian thinks this grasping can be accounted for in terms of propositional knowledge.<sup>23</sup> For instance, where  $q$  is my evidence for  $p$ , perhaps I count as grasping the connection between the evidence  $q$  and what it’s evidence for,  $p$ , provided I know the following,  $r$ :

$r$ : A belief that  $q$  justifies a belief that  $p$ .

Perhaps there are further requirements.<sup>24</sup> Regardless, if the best way to understand an internalist commitment to *Grasp* is as being committed to having propositional knowledge of an  $r$ -like proposition, then a given proposition  $P$  of the form ‘ $S$  is justified in believing  $q$  on evidence  $p$ ’ will be true only if a further proposition  $P^*$  is true where  $P^*$  is a proposition of the form ‘ $S$  knows that  $r$ ’ where  $r$  is a proposition specifying that a belief that  $q$  justifies a belief that  $p$ . Of course, according to the new relativist, the truth-conditions for  $P^*$  (of the form ‘ $S$  knows that  $r$ ’) are assessment-sensitive. We can then pose the tension like this: *if* one embraces a MacFarlane-style semantics for propositional knowledge attributions and wants to maintain this alongside an internalist account of justification according to which the ‘grasp’ condition on justification is satisfied only when the agent has propositional knowledge of a  $P^*$ -like proposition, then one can *avoid giving a relativist treatment to justification* only if willing to embrace the following unstable aporia: (i) the truth of attributions of the form ‘ $S$  is justified in believing that  $p$ ’ is *not* relative to a context of assessment; (ii) ‘ $S$  is justified in believing that  $p$ ’ is true only if it’s true that  $S$  knows a  $P^*$ -like proposition, and (iii) a  $P^*$ -like proposition gets a truth value only relative to a context of assessment.

What about attributions of *intellectual virtues*? If one is a proponent of robust virtue epistemology (of the form outlined in Chapter 7), and regards (e.g., *a la* Greco 2012) knowledge to be type identical with cognitive achievement<sup>25</sup> – *viz.*, cognitive success because of the exercise of intellectual virtues – then obviously one must be prepared to allow utterances attributing cognitive success because of intellectual virtue to get a relativist treatment. But, as I noted in the previous chapter, robust virtue epistemology

is controversial. Though we might, following Baehr (2011), be inclined to theorise about intellectual virtues entirely in the absence of any role intellectual virtues might play in the service of analysing knowledge. Baehr calls projects under this description *autonomous virtue epistemology*, and distinguishes autonomous virtue epistemology from virtue epistemology as developed in mainstream analysis of knowledge offered by Zagzebski (1996), Greco (2010) and Sosa (2007; 2009; 2011; 2015). One particular autonomous virtue-theoretic project that is increasing in interest concerns the nature of the virtue of *openmindedness*.<sup>26</sup> What does intellectually virtuous openmindedness consist in? What is its *nature*? Such projects traffic in openmindedness attributions.

*Question:* If propositional knowledge gets a relativist treatment, might openmindedness nonetheless remain 'untouched'? We've by now treaded far from the tight conceptual connections between knowledge and other notions: e.g., those knowledge stands to in relations of equivalence and reduction. But still there is reason to think a relativist treatment of 'knows' might very well bleed all the way to openmindedness, and this even if virtues such as openmindedness are, *a la* Baehr, not regarded as theoretically illuminating in an analysis of knowledge.

To see why, consider what plausibly makes openmindedness an *intellectual*, rather than, say, a moral virtue? Here's a simple framework. Any trait T is an intellectual virtue in so far as T stands relevantly connected in some way, W, with some fundamental epistemic good, E.<sup>27</sup> We can for present purposes set aside the matter of *what* the connection between the trait in question and some fundamental epistemic good, E, would look like.<sup>28</sup> What's relevant to the point I want to suggest is that, for any account of intellectual virtue that embraces this kind of model, an attractive way to fill in 'E' is with *knowledge*, that is, to say that knowledge is the fundamental epistemic good with reference to which other epistemic states, processes and traits, etc. are more or less epistemically good. Again, this is controversial. Maybe, openmindedness is intellectually virtuous in virtue of its connection toward some other end, such as true belief, or perhaps, understanding. It's beyond my present scope to take stand on this point here.<sup>29</sup> Rather, it suffices to register that knowledge is a very natural way to fill in the 'E', and it is a view that has (in different ways) enjoyed a range of support.<sup>30</sup> That said, to the extent that such a view is plausible, then relativism about knowledge attributions looks to be at tension with denying that openmindedness attributions should get a relativist treatment.

To see why, let's begin by assuming both that knowledge attributions are assessment-sensitive and further that openmindedness is

intellectually virtuous in light of some connection C holding between openmindedness and *knowledge*. Obviously, if knowledge can be very *easily* satisfied, then C holds across a wider range of cases than were knowledge very difficult to satisfy. Thus, attributions of the form 'S's inquiry I is *virtuously* openminded' will be true across a wider range of cases where knowledge is easy to satisfy than difficult to satisfy. But if the standards required for knowing are assessment sensitive, then whether C holds will be assessment sensitive. But then, the truth of attributions of the form 'S's inquiry I is *virtuously* openminded' cannot be claimed to not depend on the context of assessment. At the very least, it would be perplexing how they should not be.

Two final points of interest I want to draw attention to concern epistemic norms and epistemic value. The point about epistemic norms can be made rather simply. Many philosophers have found appealing the observation that, as Benton (2014, 1) puts it, 'knowledge provides a normative constraint or rule governing certain actions or mental states' and in particular, assertion, belief and action.<sup>31</sup> According to the *knowledge norm of assertion*, one must assert *p* only if one knows that *p*.<sup>32</sup> That knowledge is the norm of assertion is probably the most widely endorsed element of the wider knowledge-first programme. Belief, though, is (as many have thought) much like a kind of 'inner assertion', such that we can draw an analogy between belief and assertion, as inner assertion stands to outer assertion. If this is right, it's a quick step to the knowledge norm of *belief*: believe that *p* only if you know that *p*.<sup>33</sup> Furthermore, many philosophers find attractive a tight connection between knowledge and action: as Hawthorne & Stanley (2008, 578) put it, 'Treat the proposition that *p* as a reason for acting only if you know that *p*.' Obviously, if whether one knows is assessment sensitive, then the matter of whether one satisfies various knowledge rules governing mental states and actions will be as well.

Regarding epistemic value: one of the most important trends in epistemology over the past decade or so has been what Wayne Riggs (2008) has called the 'value turn' in epistemology – a shift toward much more attention to the epistemic values that frame epistemic practice.<sup>34</sup> Perhaps the most interesting upshot of the value turn has been that, following Jonathan Kvanvig (2003), many philosophers have come to regard the following as a condition of adequacy on any analysis of knowledge: that it not be logically incompatible with a plausible datum about the *value* of knowledge, which is that the epistemic value of knowledge exceeds that of merely justified true belief that falls short of knowledge.<sup>35</sup> Famously, as Kvanvig (2003) argues, reliabilist accounts of knowledge according

to which knowledge is reliably produced true belief fail to satisfy this condition of adequacy. Reliability, he argues, is an epistemically valuable property of a belief to have in so far as it renders the belief likely to be true. But then, it's hard to see why this properly should add any *additional* epistemic value to an *already* true belief. Thus, as the argument goes, the epistemic value of reliability is 'swamped' by the epistemic value of truth in a way that leaves the view that knowledge is reliable true belief with no way, in principle, to preserve the insight that knowledge is epistemically more valuable than justified true belief that falls short of knowledge. Reliabilists have attempted to defend their view against this objection,<sup>36</sup> and moreover, other views have recognized the 'swamping problem'<sup>37</sup> as a problem that must be avoided in order for an analysis of knowledge to meet the value-driven adequacy condition on knowledge.<sup>38</sup>

Let's now connect this strand of thinking in epistemology with the thought that whether 'S knows that  $p$ ' is true is assessment sensitive. In particular, if knowledge attributions get a relativist treatment, then where does this leave claims of epistemic value? The answer, in short, is: not insulated in a protective non-relative bubble. Let's operate with the assumption that knowledge is, in fact, more valuable than mere true belief, such that: a given token of knowledge – Tim's knowledge that  $W$  is the way to Larissa' is more epistemically valuable than a corresponding *mere true belief* token with the same content, that  $W$  is the way to Larissa. Call the proposition that  $W$  is the way to Larissa 'L'.

Now, suppose Tim actually does have a belief, which is true, that L. Call this state of affairs  $\alpha$ . Let  $\beta$  be a possible state of affairs that is just like  $\alpha$  except that in  $\beta$  Tim's belief that L is *merely* true, and not knowledge.

*Value Question:* is  $\alpha$  a more epistemically valuable state of affairs than  $\beta$ ?

If we are assuming that knowledge of a given proposition  $p$  is more epistemically valuable than a corresponding mere true belief token that  $p$ , then to answer the value question, it *matters* whether  $\alpha$  is a state of affairs where *knowledge* is present. Suppose now that Liz utters ' $\alpha$  a more epistemically valuable state of affairs than  $\beta$ '. The context in which she asserts this,  $c_1$ , is the context of Tim. Now, suppose that Charles evaluates the truth of Liz's assertion from context  $c_2$ . Working with MacFarlane's (2014) view, we'll say that Liz's assertion, as uttered by her, and as assessed from  $c_2$ , is true only if  $\alpha$  is a state of affairs in which Tim can rule out alternatives relevant in Liz's context (which is, again, Tim's). Stipulate that this is an

'easy' context, suppose, where Tim counts as knowing L provided Tim can rule out that L is not the way to any of the other salient communities in Thessaly such as Giannouli, Platykampos or Nikaia, something he plausibly easily counts as doing in  $c_1$  given that he's seen the *sign* on which is printed 'Larissa' pointing in the direction of  $W$ , and also further signs indicating that the roads to Giannouli, Platykampos and Nikaia fork off in different directions from  $W$ . Charles, however, thinks to himself 'might the sign to Larissa have recently been switched in an extremely skilled and stealthy fashion by Meno, so that it is now misleading in a way that would not be detected by human vision?' On MacFarlane's view, the claim that Tim knows L would not be true, as assessed by Charles, given that Tim cannot rule out the alternatives relevant in Charles's context, which includes the alternative that Meno stealthily switched the sign in a way that would evade detection. Now, again, if we are assuming that knowledge of a given proposition  $p$  is more epistemically valuable than a corresponding mere belief token that  $p$ , whether Liz's assertion that  $\alpha$  is a more epistemically valuable state of affairs than  $\beta$  is true, as evaluated by Charles, will depend on whether  $\alpha$  is a state of affairs where knowledge is present, as evaluated by Charles. As we've just said, though,  $\alpha$  is not a state of affairs where knowledge is present, as evaluated by Charles. And so Liz's assertion that  $\alpha$  is a more epistemically valuable state of affairs than  $\beta$  comes false as asserted by Liz at  $c_1$  and as assessed by Charles at  $c_2$ , even though it comes out true as assessed by Liz at  $c_1$ . Thus, if knowledge attributions get a relativist treatment, then at least on one standard way of thinking about epistemic value in the contemporary literature, answers to the value question can end up true relative to some contexts of assessment, false relative to others.

Let's now take stock. The Epistemological Ramification Question (ERQ) asked: if 'knows' gets relativist treatment, then since knowledge relates intimately with other epistemic concepts, do any other epistemic concepts also need a relativist treatment? The answer, as I've developed it in this section, is clearly 'yes', though as we've seen the *extent* of the bleedover from relativism about 'knows' to other epistemic notions depends, in various ways, on which epistemological views we already embrace. As the table below shows, a truth-relativist semantics for knowledge attributions will force (in light of the relevant commitments noted in the left-hand column of the table) a relativist treatment of at least the following: evidence, knowledge-how, understanding-why, justification, norms governing assertion, belief and action, intellectual virtues and epistemic values.

---

E = K (e.g. Williamson)	Evidence
Reductive intellectualism about knowledge-how (e.g. Stanley)	Knowledge-how
The ‘traditional view of understanding’	Understanding why
(i) Traditional JTB analyses; (ii) certain kinds of internalists	Justification
Knowledge norm of assertion, knowledge norm of belief, knowledge norm of action	Norms governing assertion, belief and action
(i) RVE; (ii) K-aimed (teleological) accounts of intellectual virtues	Intellectual virtues
$V_e(K) > V_e(TB)$ (e.g., Meno Assumption)	Epistemic value

---

The above table is, I think, very much incomplete. One needn't be a card-carrying 'knowledge-firster' to see that whether a range of claims attributing core epistemic notions are true depends (in different ways) on what one knows. This is much weaker than the characteristic K-first claim that other epistemic notions are analysable *in terms of* a more fundamental notion: knowledge. But even with this much less controversial idea in play, it's easy to see how from a truth-relativist semantics for knowledge attributions, we're on a quick path to something like a *wholesale epistemic relativism*: that is, relativism about many or perhaps even most of the notions epistemologists study, along with values and norms that structure our practice of attributing and analyzing these notions. And this is just what you'd expect, given that most epistemologists will tell you that epistemology is, in short, the theory of knowledge.<sup>39</sup>

### 8.3 A parity argument: non-factivity

I take the previous section to have established that the following is highly plausible: *if* propositional knowledge gets a relativist treatment, then so will a range of other notions important in epistemology. In short, we can't very well give 'knows' a relativist treatment in epistemic isolation and carry on doing epistemology just the same as before.

Thus, in the wake of a relativist treatment for knowledge attributions plausibly follows a kind of 'wholesale' epistemic relativism – that is, a need to recognise the truth of claims featuring epistemic notions intimately related to knowledge as *themselves* depending on a context of assessment

provided we have it as collateral that the matter of whether one counts as knowing a proposition depends on a context of assessment.

However, I now want to carefully qualify the conditional I've just defended. It really only goes through if we accept already a certain, admittedly very standard (especially since the early 1990s)<sup>40</sup> background assumption, one about the relationship between the linguistic theory of knowledge attributions and the epistemological theory of knowledge. Call this the *illumination thesis*:

*Illumination Thesis*: the linguistic theory of knowledge attributions should guide/inform/illuminate the epistemologist's theory of knowledge.

Obviously, if the illumination thesis is simply *false*, then facts about how ordinary speakers use 'knows' – *viz.*, facts about how 'knows' is *used* which are indirect evidence for what the meaning of 'knows' is – would be irrelevant to how *epistemologists* should theorise about knowledge.<sup>41</sup> Though most people *don't* think it's false. The standard story is roughly this. Theorising about the knowledge on the basis of how 'knows' is used is legitimate because: (i) how we use 'knows' is a guide (modulo pragmatic and speaker error explanations of our use data) to the truth conditions of knowledge attributions and thus the meaning or definition of 'knows' – *viz.*, necessary and sufficient conditions for its correct *application*;<sup>42</sup> and (ii) the epistemologist is for her part looking for something very similar to a definition for knowledge in attempting to specify necessary and sufficient conditions in the service of analysing knowledge.

But this standard story behind the illumination thesis is itself premised upon an even *more* fundamental assumption: that the concept of 'knows' as used by ordinary speakers is in fact the same, or at least very close to, the one epistemologists are interested in. If *this* assumption is false, then we should resist, rather than embrace, the illumination thesis.

Allan Hazlett (2010) has recently argued that the illumination thesis is in fact false, and he's done this precisely by arguing that the concept of 'knows' as used by ordinary speakers *cannot* be the one that epistemologists are theorizing about under the description of 'knowledge.' I want to now consider how Hazlett's argument, which is creative, is supposed to work, and then to put a few pieces together: that is, to show how Hazlett's rationale, with a few refinements, might be combined with the results of the previous section to support a new strategy epistemologists

might have for rationally dismissing what we've called new epistemic relativism, the epistemic relativism of Chapter 7 – and it is a strategy which bypasses entirely the Achilles' Heel Question and the 'semantic stalemate' that it threatens.

### 8.3.1 *Hazlett's anti-illumination argument*

Following here Turri (2011, 144–5), we can present the key steps in Hazlett's argument as follows:

#### *Hazlett's Anti-Illumination Argument*

1. Any non-factive concept of knowledge is epistemologically uninteresting. (Premise)
2. The ordinary concept of knowledge is non-factive. (Premise)
3. So the ordinary concept of knowledge is epistemologically uninteresting. (From 1 and 2)
4. If the ordinary concept of knowledge is epistemologically uninteresting, then ordinary knowledge ascriptions are epistemologically uninteresting. (Premise)
5. So ordinary knowledge ascriptions are epistemologically uninteresting. (From 4 and 5)

The conclusion, (5) is tantamount to the denial of the illumination thesis. The perspective Hazlett (2010, 500; 522) takes on his argument is as follows:

I'm suggesting [...] a divorce for the linguistic theory of knowledge attributions and traditional epistemology [...] Traditional epistemology and ordinary language epistemology (as we might call the theory of knowledge attributions) would both be best served by going their separate ways.

As we said, the illumination thesis is undermined if the ordinary concept of knowledge and the one epistemologists are interested in is not the same one. The crux of Hazlett's argument is to suggest exactly this, by pointing to a property, factivity, that he insists is both essential to the epistemologist's conception of knowledge, and which is *not* a mark of the ordinary folk concept 'knows'.

Obviously, (2) is the most controversial premise. And indeed, most of Hazlett's paper is aimed at establishing it.<sup>43</sup> I'll not weigh in on the matter of whether Hazlett's (2) is in fact true, as this is irrelevant to the point I'll be making.<sup>44</sup> Rather, what's of interest for the present purposes

is this: provided (1) is true – that is, that provided that the *epistemological concept of knowledge is a factive one*, then it looks like the *more* evidence we have for thinking (2) is true (i.e. that the ordinary concept of knowledge is non-factive), the *less* reason we think have for thinking that the illumination thesis is true – *viz.*, that the linguistic theory of knowledge attributions should guide/inform/illuminate the epistemologist's theory of knowledge. And, more generally: to the extent that use data supports regarding our ordinary concept of knowledge having *any given property P* which is not shared by the epistemologist's conception of knowledge, we have reason to regard the illumination thesis as false.

Premise (1) of Hazlett's argument relies on the thought that factivity is *essential* to the concept of knowledge as studied by epistemologists, such that, a non-factive concept of knowledge is not going to be epistemologically interesting.<sup>45</sup> And he's surely right about this.<sup>46</sup> But what *else* is essential to the epistemologist's conception of knowledge?

Here's an answer I want to now defend. It's essential to the epistemologists's conception of knowledge that it can support and sustain the practice of *epistemological discourse*. Let's say a given concept of knowledge can support the practice of epistemological discourse only if, the debates that are central in epistemology continue to *make sense* were we to imagine that epistemology is centred around that concept so conceived; that is, debates that are typical of first-order epistemological theory could continue to take at least roughly the same kinds of shapes they presently take, in practice.

One reason that the epistemologist's conception of knowledge cannot be non-factive is that a non-factive concept of knowledge simply can't support epistemological discourse in the sense just described (the case of the Moorean versus the sceptic from Chapter 1 being an illustrative example). Debates would not even take roughly the same shape were these debates about *non-factive-knowledge*.<sup>47</sup>

In Chapter 1, I argued that mainstream epistemology takes for granted more than just factivity, but further, something else: a commitment to regarding epistemic facts as objective in character. This was a claim that took some defending, but the way I argued for it was, in short, *via* two key premises. I argued that (i) if most disagreements at the first-order in epistemology are ones where both sides to the disputes are disposed to behave, in their use of language, as if they believe something *p* to be common ground in the context of their dispute, then *p* is a *meta-epistemological* commitment of both sides of the first-order dispute. And I argued further for the descriptive claim that (ii) most first-order disagreements in epistemology (i.e. consider the debate between

Goldman and Feldman about the nature of epistemic justification) are ones where both sides to the disputes are disposed to behave, in their use of language, as if they believe a commitment to objective epistemic facts is common ground in the context of their dispute. And so I concluded that a meta-epistemological commitment of most first order-disagreements in epistemology is a commitment to objective epistemic facts. This is of course not to say that what is presupposed by most first-order debates in epistemology is *right*. (After all, maybe there aren't any objective epistemic facts, and so what's pragmatically presupposed by most first-order debates is just wrong.) But it remains that if my arguments from Chapter 1 hold, a revealed meta-epistemological commitment in mainstream epistemology is a commitment to epistemic facts with an objective profile, a commitment which, were it not a pragmatic presupposition in the common ground of most first-order debates, we'd not be able to explain the characteristic *shapes* such debates actually take in practice.

That being said, I now want to motivate a Hazlett-style dilemma to the proponent of assessment-sensitive knowledge attributions. The dilemma is that the more persuasively MacFarlane can support his thesis – that the ordinary concept of knowledge is to be given a relativist treatment – the stronger the case becomes for *rejecting* the illumination thesis, *viz.*, the thesis that the linguistic theory of knowledge attributions should guide/inform/illuminate the epistemologist's theory of knowledge. And, as I suggested, without the illumination thesis in hand, we have no reason to take MacFarlane's variety of epistemic relativism as one that should have relevance to epistemological theory. Call this dilemma the *anti-illumination dilemma*, a dilemma we can model off Hazlett's non-factivity argument for the same conclusion.

#### *Anti-illumination dilemma*

6. Any relativist concept of knowledge is epistemologically uninteresting. (Premise)
7. The ordinary concept of knowledge is relativist. (From MacFarlane's master argument, Chapter 7)
8. So the ordinary concept of knowledge is epistemologically uninteresting. (From 1 and 2)
9. If the ordinary concept of knowledge is epistemologically uninteresting, then ordinary knowledge ascriptions are epistemologically uninteresting. (Premise)
10. So ordinary knowledge ascriptions are epistemologically uninteresting. (From 4 and 5)

Premise (10) is tantamount to a denial of the illumination thesis. If (10) is true, then it looks like *if MacFarlane is right* that 'knows' is assessment-sensitive, then the ordinary concept of knowledge is *epistemologically* uninteresting, and moreover, so are ordinary knowledge ascriptions. In this respect, the anti-illumination dilemma proposes that the more convincing MacFarlane is that the ordinary concept of knowledge is assessment-sensitive, the less *epistemologists* should be concerned by this result.

A defence for (6) in the argument is offered by combining two ideas. The first idea is the one – just outlined – that a metaepistemological commitment of most first-order debates in epistemology is a commitment to objective epistemic facts. The second idea is the conditional claim defended in the previous section: the conditional claim was that *if* the epistemological concept of knowledge were given a relativist treatment (a la MacFarlane), then an implication would be a kind of *wholesale* epistemic relativism, one that is incompatible with the metaepistemological commitment revealed by most first-order projects in epistemology to objective epistemic facts – a descriptive claim I've argued to be true in Chapter 1 by linking metaepistemological commitments with pragmatic presuppositions and then outlining what we actually find to be pragmatically presupposed, in the common ground, of most first-order disputes. Thus, just as Hazlett defends (1) of his argument by drawing attention to how mainstream epistemology is premised upon a factive concept of knowledge, in such a way that a non-factive concept will be epistemologically uninteresting, I've suggested that we accept (6) for parallel reasons.

Whereas Hazlett actually defended his own premise (2) – that the ordinary concept of knowledge is non-factive – it is *MacFarlane* who is aiming to establish premise (7) of the anti-illumination dilemma: that the ordinary concept of knowledge is in fact relativist, *viz.*, that 'knows' as used in ordinary language is assessment-sensitive.

And it is here where I want to emphasise the deep dilemma for MacFarlane. In the previous chapter, I offered epistemologically oriented considerations for thinking that MacFarlane's own assessment-sensitive semantics for 'knows' has epistemological costs. In the first section of this chapter, however, I backtracked and conceded that I'm in no position to insist that the costs I've claimed his view incurs are obviously *worse* than the costs incurred by the other leading competitor semantics for knowledge attributions. And so MacFarlane's view was granted (at least, in light of what was established in Chapter 7) a 'semantic stalemate', a concession according to which the burden remains on the relativist to

show why the problems raised in the previous chapter are *less* costly than the problems commonly associated with the other competitors.

And now, as I've attempted to set things up, the dilemma facing MacFarlane has been repositioned as follows: the more compelling an argument MacFarlane can brandish to the effect that 'knows' should get a relativist treatment – that is, the more reason we have to accept premise (7) of the anti-illumination dilemma – the more reason the epistemologist has for setting aside new relativism on the grounds that the concept of knowledge of ordinary language, the one the relativist aims to show is assessment-sensitive, is epistemologically uninteresting.

#### 8.4 Concluding remarks

I want to conclude by briefly resituating the dilemma raised at the end of this chapter. MacFarlane's argument for an assessment-sensitive semantics for 'knows' just *is* an argument for (7), that the ordinary concept of knowledge is a relativist (e.g. assessment-sensitive) one. Now here's the Sophie's Choice:

- (i) If MacFarlane makes the argument for (7) poorly, that is, if we have no good reason to embrace a truth-relativist semantics for knowledge attributions *over* competing semantics for knowledge attributions (e.g. insensitive invariantism, SSI, contextualism) then 'new epistemic relativism' is a view that mainstream epistemologists can rationally dismiss on the grounds that any semantics for knowledge attributions should be interesting to epistemologists only if it fares at least as well as competitors.
- (ii) *However*, if MacFarlane doesn't make the argument for (7) poorly, but instead, if he makes it *very well*, so that we have compelling reason to think that the ordinary concept of knowledge is a relativist one as MacFarlane tells us, then to the extent that (6) in the anti-illumination argument is correct, then epistemologists are going to be entitled just as well to proceed as they have before way as before, given that (6) and (7) imply that the truth of MacFarlane's thesis, qua a thesis about knowledge attributions featuring the ordinary concept of knowledge, is not an epistemologically interesting one.

Either way, it looks like new epistemic relativism should lack import in epistemology.<sup>48</sup>

# 9

## Metaepistemology and Relativism

*Abstract.* This chapter has two primary aims. Firstly, the conclusions drawn from Chapters 2–8 will be situated within a wider context: that of the complicated relationship between metaepistemology, realism and relativism, with an eye to answering some questions left outstanding at the end of Chapter 1. The second thing I want to do is to motivate, in light of what's been argued for already, an entirely different way of thinking about the relationship between mainstream epistemology and the kind of challenge to it that arguments for epistemic relativism stand to pose. If I am right, at least some arguments for epistemic relativism (particularly, new epistemic relativism) *can*, and despite what was insinuated at the end of Chapter 8, *do* have an important kind of relevance to mainstream epistemology, even if more traditional arguments of the sort canvassed in Chapters 3–5 do not. Though the *kind* of relevance is – I'll argue – of a very different sort than is ordinarily thought.

### 9.1 The lay of the land

I now want to situate some of the overarching conclusions drawn from Chapters 2–8 within the wider narrative with which I began. One such conclusion was that traditional arguments for epistemic relativism – e.g. argument styles which we can locate within the Pyrrhonian, Dialogic and Incommensurability-style template argument forms – were ultimately not compelling, and so failed to motivate a relativistic picture of epistemic facts. The common factor seemed to be that traditional argument strategies failed to ultimately give us cause to prefer relativism to scepticism (as well as to, in the case of arguments which followed the incommensurability–circularity–relativism

sequence, other non-sceptical alternatives<sup>1</sup>). However, ‘new epistemic relativism’ – the kind of view championed on semantic grounds by John MacFarlane, turned out to be a different story altogether, one that we saw simply cannot be dismissed on the basis of the familiar kinds of considerations that plagued more traditional-styles of arguments for epistemic relativism.

After considering arguments for and against New Epistemic Relativism, the conclusion reached was a precarious one: that *to the extent* that this view successfully motivates a relativist treatment of ‘knows’ (something I raised some epistemologically oriented doubts about in Chapter 7), its doing so should be best thought of as threatening not relativism *in epistemology*, as we’d be initially inclined to think after examining the anticipated ‘fallout’ (in epistemology) of accepting a MacFarlane-style semantics for ‘knows’,<sup>2</sup> but more accurately, something akin to a Hazlett-style *divorce* for the theory of knowledge attributions (which concern the ordinary concept of knowledge), and the *epistemological* concept of knowledge – the one epistemologists are in fact interested in. I’ll return to this issue. But the broader conclusion reached over the course of Chapters 2–8 – for ease of reference, in what follows, call this the *primary conclusion* – is that that *neither* traditional forms of epistemic relativism, nor new epistemic relativism as motivated on semantic grounds ultimately, at the end of the day, poses a credible ‘threat’ to mainstream epistemology, in the sense that neither ultimately gives us a good reason to think that the kind of mainstream meta-epistemological realism that, in Chapter 1, revealed itself as firmly in the background of most first-order projects in mainstream epistemology, should be supplanted with a picture on which the epistemological facts epistemologists take themselves to be trafficking in when pursuing first-order projects are to be understood as relative ones.<sup>3</sup>

In this chapter, I want to put this primary conclusion to work in the service of bringing this book to a close, in particular, by doing two final things:

- (i) Firstly, I want to situate the primary conclusion in wider context, by showing how this conclusion (along with some further premises) helps to motivate an answer to the *incompatibility question*, which was left open at the end of Chapter 1. This question was, roughly, whether epistemic relativism is *fundamentally* incompatible with mainstream epistemology. The answer I’ll opt for here is an affirmative one, though to appreciate why, we’ll need to do a bit of stage setting.

- (ii) The second thing I want to do is to motivate, in light of what's been argued for already, an entirely different way of thinking about the relationship between mainstream epistemology and the kind of challenge to it that arguments for epistemic relativism stand to pose. If I am right, at least some arguments for epistemic relativism (particularly, new epistemic relativism) *can* have an important kind of relevance to mainstream epistemology, though the *kind* of relevance is – I'll argue – of a very different sort than is ordinarily thought.

## 9.2 The incompatibility question

I'm going to take it for granted from here on out that the primary conclusion I've argued for is correct. And if it is, then this goes some way to vindicating as not especially problematic the observation we began with in Chapter 1: that first-order projects in mainstream epistemology typically proceed as though epistemic relativism is false or at least can be set aside. However, as I indicated toward the end of Chapter 1, there might be a much deeper explanation here. Maybe epistemic relativism is just plain *incompatible* with the assumptions that first-order projects in epistemology have to take for granted to get their debates off the ground. If such an incompatibility can be demonstrated, then it's not merely that (typical) first-order projects in mainstream epistemology have no philosophically motivated reason to amend or abjure their positions in light of the arguments epistemic relativists have been able to offer; it's that typical first-order projects in mainstream epistemology *can't* do so – because these very projects *depend* on metaepistemological commitments that are incompatible with epistemic relativism.

This is the question of incompatibility: and providing an answer to it involves, as I suggested in Chapter 1, doing some careful thinking about the (very) complex relationship between *metaepistemological realism* (a commitment to which, to emphasise, I argued in Chapter 1 is to be found in the common ground of most first-order projects in mainstream epistemology<sup>4</sup>) and relativism. I want to now revisit these complexities, and in doing so, make it clear why we can – now with our primary conclusion in hand – answer the incompatibility question in the affirmative.

Metaepistemological *realism*, as I defined it in Chapter 1, involves two central components: an *existence* component and an *independence* component.<sup>5</sup> By the existence component of metaepistemological realism, the metaepistemological realist (hereafter, the realist) is committed<sup>6</sup> to the existence of epistemic facts, something that the metaepistemological

error-theorist and the non-cognitivist deny. And, further, the realist must also regard the epistemic facts countenanced to be suitably *mind-independent*.

Being committed to *trivial* mind-dependencies is of course compatible with satisfying the ‘independence leg’ of realism. For example, Williamson (2000), in claiming that knowledge is a mental state, is committed to the view that knowledge-facts are, in a trivial an uninteresting sense, mind-dependent. But obviously Williamson isn’t an anti-realist about knowledge in virtue of what he is committed to here.

Thus, as it was suggested in Chapter 1, denying the ‘independence leg’ of meta-epistemological realism is going to involve a commitment to regarding epistemic facts – *viz.*, states of affairs consisting in the instantiation of commonly conceived epistemic properties in the world – to the effect that their obtaining is *non-trivially* mind-dependent; that is, that epistemic facts are dependent on our mental lives in ways that (put roughly) go beyond what is necessary for there to be such facts at all.<sup>7</sup>

Traditionally at least, *idealism* – the unpopular thesis according to which reality ‘consists exclusively of minds and their ideas’<sup>8</sup> – has been regarded (along with Dummett’s (1978) semantically driven anti-realism and Goodman’s (1975; 1978) brand of constructivism<sup>9</sup>) as among the most obvious strategies by which the *independence* leg of realism is opposed.<sup>10</sup> However, things are in fact very thorny in this area, and I want to now revisit this issue, as it was first raised in Chapter 1, and to outline carefully just how it connects with the matter of the relationship between metaepistemological realism and epistemic relativism (and specifically, as regards what would have to be the case were the latter to count as falling *within* rather than *outside* the camp of the former).

Here, an example will be helpful, one which generates a puzzle, my answer to which reveals what I think is ultimately at issue when answering the compatibility question. Recall again the toy view entertained briefly in Chapter 1 – a view that, as far as I’m aware, no one writing today endorses, which we can call *Berkeleyan idealism about epistemological facts*. On this imagined view, it is claimed that there *are* epistemological facts,<sup>11</sup> and further, that reality *itself* is mental in nature, and so the states of affairs that obtain when one knows a given proposition will necessarily be mental, in the following sense: that *anything we know* we know because some mental state of affairs obtains. This is some wild supposing, but let’s continue with it. The puzzle emerges when the Berkeleyan metaepistemological idealist (for convenience, call this individual ‘George’) surprises us by telling us that we are simply wrong to regard him as an *anti*-realist. George tries to tell us that he is a

full-blooded *realist*. (George would be wrong of course, but getting clear about exactly why would be instructive.)

His story about why we should think of him as a realist goes like this: firstly, he tries to distance himself from the error-theorist by explicitly saying he embraces the existence of, and does not deny, *epistemic facts*, and on this basis, he tells us we should agree that the *existence* component of (metaepistemological) realism is secured on his view. Further, to convince us he doesn't deny the independence leg (a leg which we are, previously, quite sure he does deny) George says, describing the Berkeleyan metaepistemological idealist's view, that: 'Given that the world is just an idea in the mind (and thus, given that for instance the reality we aspire to represent is itself of a mental constitution), it's from my perspective going to be entirely *trivial* that epistemic facts are mind-dependent. But we've already agreed that one does not run against the independence leg of realism simply in virtue of allowing that epistemic facts are *trivially* mind-dependent'.

Now I think we need to be careful in replying to this imagined piece of reasoning. After all, *if* the Berkeleyan conception of what an epistemic fact is – one replete with the Berkeleyan's material idealism about the world which beliefs aspire to represent – is a *correct* conception, then an epistemic fact (for example, 'S knows that *p*') will be a state of affairs consisting in an individual, who is an idea in the mind, instantiating a property, e.g., knowledge, which the individual who is an idea in the mind instantiates in virtue of having a successful representation (whatever this involves) of some other state of affairs *that is also an idea in the mind*. The *mind-dependence* of epistemic facts, if Berkeleyan empirical idealism is true, would turn out to be a *trivial* dependency. But of course this is a perplexing result: the thought that the Berkeleyan idealist should ever have any sort of claim to make that she is a metaepistemological realist seems to water down realism beyond recognition.

The way out of this perplexity, which I alluded to in Chapter 1, and which I think best sets the scene for how to think about the relationship between metaepistemological realism and relativism (such that we'll have a clearer idea about whether the latter is compatible or incompatible with the former), was to begin by pairing together two connected insights.

The first insight is this: that there are *some ways* that a view can opt for the mind-independence of a class of fact in question which will have, as an unavoidable consequence, that the *existence* leg of realism is also *de facto* denied, and this is so *even if* the view *claims* to countenance the existence of the relevant class of facts; that is, on at least some kinds of

views, claims about the *way* facts which the view claims to countenance are regarded as mind-independent can have, as a knock-on effect, that the *existence* leg of realism simply can't count as upheld. This general schema is thus false: that for any (would-be) realism, whether the existence component is denied always remains open provided independence is denied.

The second insight has to do with how this *works*. And here our toy Berkeleyan meta-epistemological idealist's view will be illustrative – *viz.*, of just how *the way* the facts one claims to countenance are mind-dependent can matter for whether the existence leg of realism also has to be given up. The idea proceeds like this: following Cuneo (2007), we should insist that – and this is the second key insight – *not everything gets to count as an epistemic fact*; there are *conceptual limits* to what counts as an epistemic fact. To satisfy the existence leg of realism, you must countenance not only something you refer to as an 'epistemic fact' (e.g. something which, despite what you say, might actually be – and unbeknownst to you – a moral fact, or a horse fact), but something that falls *within these conceptual limits* and not outside of them.

Now, pairing these two insights together: it should be clear how, if one claims to embrace the existence of epistemic facts but in doing so opts for a characterisation of these facts on which epistemic facthood is, say, a radically mind-dependent affair, then *what one's got left isn't necessarily going to be epistemic facts anymore*. It might be something else, contrary to what this individual is telling us. Accordingly, what I think we should be asking ourselves – for the purpose of thinking about whether epistemic relativism is even compatible with metaepistemological realism, or whether it is rather a form of anti-realism – is whether the kind of epistemic facts embraced by the epistemic relativist fall *within plausibly circumscribed conceptual limits* of epistemic facthood. If not, then the epistemic relativist is for better or worse an anti-realist regardless of her own protestations.

Having set things up this way, the next relevant point will be, of course, what *determines* what falls within the conceptual limits of an epistemic fact? The story here had better be a good one. And as I noted in Chapter 1, I think Cuneo's (2007) story here is very reasonable;<sup>12</sup> on this view, *epistemic facts*, as *such*, must satisfy certain common-sense *platitudes*, as Cuneo puts it, platitudes about the (i) content and (ii) authority of epistemic facts. Facts that don't satisfy these platitudes aren't epistemic facts.

Regarding *content* platitudes: Cuneo's claim was that, just as content platitudes *vis-à-vis* moral facts 'congregate around the notion of human

well-being' in such a way that if a fact is entirely orthogonal to human well-being, we are wont to call it a *moral* fact (p. 58), likewise, 'content platitudes with respect to epistemic facts cluster around the notion of *accurate representation*' that is, of '*representing reality aright*'.<sup>13</sup>

But epistemic facts, as Cuneo remarks, must also satisfy commonsense platitudes about the *authority* of epistemic facts. Not just any fact that satisfies the content platitude is authoritative in the way that epistemic facts are. Here, recall, the line was that:

[Epistemic facts] are *prescriptive* ... they are, imply, or indicate reasons for properly situated individuals to behave in certain ways ... *regardless* of whether these agents care about conducting their behaviour in a rational way, whether they belong to a social group of a certain kind, or whether they have entered into a social agreements with others. (*Ibid.*, p. 59)

That epistemic facts, of the sort the realist countenances, are prescriptive in this sense – *viz.*, that their reason-giving force isn't just a matter of what we care about – is also a point that is granted by individuals who unlike Cuneo distinguish themselves as *denying* the existence leg of realism. Specifically, what card-carrying epistemic error-theorists, such as Jonas Olson (2009), want to deny, *qua epistemic error theorist*, is that there are epistemic facts with exactly this kind of a prescriptive profile. Olson in fact takes care to point out that when the error theorist purports to deny the existence of epistemic reasons, what she is denying is *not* that there are (for instance) goal-dependent or role-dependent reasons to manage one's cognitive life in certain ways<sup>14</sup> – the epistemic error theorist, Olson emphasizes, can unproblematically accept the existence of these kinds of reasons<sup>15</sup> – but rather the error theorist denies that 'there are reasons for agents to believe ... [which do] ... *not* depend entirely on agents' desires, ends, roles, or activities.' Thus, just as the authority platitude characterises the kind of epistemic facts the error theorist self-confessedly wants to deny under the description of an epistemic fact, the authority platitude captures exactly the kind of thing Cuneo tells us you have to *countenance* if you're claiming (*unlike* the error theorist) that you do recognise *epistemic* facts.<sup>16</sup> And the key thrust of the authority platitude also enjoys explicit expression in discussions more generally of differences between the epistemic and practical.<sup>17</sup>

Now, with reference to content and authority platitudes which constrain epistemic facthood, it's relatively straightforward to see why the toy Berkeleyan view in fact *fails* to countenance the existence of

*epistemic facts proper*: the kind of entity the Berkeleyan regards as an epistemic fact is going to fall outside the conceptual limits of epistemic facthood precisely *because* such an entity will not be one for which the authority platitude could very plausibly, or perhaps even coherently, apply.<sup>18</sup> We can illustrate this with a simple example: stipulate that the Berkeleyan metaepistemological idealist's picture of 'epistemic' facts is correct. Stipulate now further that Goldman 'Berkeley-knows' some epistemic fact, *B*. In virtue of what could *B* imply or indicate any epistemic reasons for Feldman? It would at least be an utter mystery how this should be so.<sup>19</sup> The Berkeleyan, despite whatever protestations she might have, doesn't countenance *epistemic facts* when claiming that she does.

*Question*: Does the epistemic relativist do better than the Berkeleyan does in preserving epistemic facts? Put differently, is the epistemic relativist in the same boat with the Berkeleyan in the following sense: if the epistemic relativist tells us that she countenances the existence of epistemic facts, do the facts she countenances fall *within* plausible conceptual limits of epistemic facthood? Getting to the crux of things: *Are relative epistemic facts such that the content and authority platitudes apply to them?* It's plausible that the content platitude does, and I won't question this point.<sup>20</sup> What's (much) more interesting is the authority platitude.

If I am right, how we can provide a principled answer to the compatibility question rides, very much, on this very issue. Though here there was a 'catch' – something of a cliffhanger, with which we ended Chapter 1, and it's a cliffhanger that, armed with our primary conclusion, we are now in a position to resolve. The cliffhanger was, exactly, this: *if* the Cuneo-style (non-relativist) realist picture of epistemic facthood is correct, then *what it takes* for a metaepistemological theory *T* to respect the authority platitude is going to take a particular shape, where the *authority* that epistemic facts exhibit will involve a kind of *inter-perspectival categoricity*, where the prescriptive force of these facts transcends *not only one's desires and preferences, but also what epistemic frameworks/perspectives one subscribes to*, or what 'epistemic culture' one belongs to. For example, if realists like Cuneo or Timothy Williamson are correct, then the authority of facts of the form 'evidence *E* justifies belief *B*' will have prescriptive force for *S*, regardless of not only *S*'s adopted aims, but *also* regardless of whether *S* is, say, a Zande who wishes to abstain from judgment on the matter of *B* until the poison oracle's verdict is in.

*However*, and this is the flip side of the coin: *if* either traditional or new epistemic relativism is *correct*, then *what it takes* to satisfy the authority platitude constraining epistemic facthood will take on a very

*different* shape, one where the authority of epistemic facts will be understood as having *intra-* rather than *inter-perspectival categoricity*, such that the prescriptive force of epistemic facts, as such, *does* transcend one's desires and preferences (unlike the kind of hypothetical reasons the error theorist like Olson can accept), but *not one's epistemic frameworks/perspectives*. For instance, in the case of traditional (e.g. Rorty-style) epistemic relativism, the epistemic relativist will view Galileo's belief in Copernicanism to be justified relative to Galileo's scientific grid, and the fact that his belief attains this kind of relative justification is regarded by Rorty as authoritative albeit *within though not outside* that particular grid. Likewise, on MacFarlane's model, epistemic facts (e.g. attributions of knowledge to a subject) can be upheld as having an authority that does transcend the agents' desires and goals (after all, desires and goals aren't afforded semantically significant role), but the authority's categoricity will be *intra-perspectival*, where the epistemic 'perspective' is keyed to the context of assessment.

And now for the two claims that frame, explicitly, what was the cliff-hanger at the end of Chapter 1:

- (i) *If* the Cuneo-style (non-relative) realist is right and the relativist wrong, then it looks like the epistemic relativist is *not* a metaepistemological realist, and this because the epistemic relativist fails to satisfy the Cuneo-style-realist's idea of what it takes to satisfy the authority platitude according to which the notion of authority should be unpacked *inter-perspectivally*. That is: if we have more reason to embrace a Cuneo-style (non-relative) realism about epistemic facts than we do relativism, then we also have a reason to insist on a certain conception of what is required to satisfy the authority platitude, one according to which *intra-perspectival* authority (of the sort that characterize what would be the epistemic relativist's claimed epistemic facts) is not sufficient.
- (ii) *However*, the other side of the coin is that, if we have good reason to think the epistemic relativist is *right* – i.e. if arguments for epistemic relativism are compelling ones – then it looks like the epistemic relativist *is* going to be best classed as a *metaepistemological realist*, and this because the epistemic relativist does satisfy the *authority platitude* unpacked with an *intra-perspectival* gloss as it must be if epistemic relativism is correct – *viz.*, that epistemic facts have an authority that transcends desires and goals but *not* epistemic frameworks/perspectives.<sup>21</sup> And so: if we have compelling reason to embrace the epistemic *relativist's* picture of epistemic facts, then

we have *as well* reason to insist on a certain conception of what is required to satisfy the authority platitude, one according to which *intra-perspectival* authority is sufficient.

Bearing this dilemma in mind, we can now put everything together: I said that, armed with the primary conclusion, we could defend an affirmative answer to the compatibility question, and things are suitably set now to see why. With reference to the primary conclusion that's been established, *it's not the case that* we have any compelling reason to accept the picture of epistemic facthood the epistemic relativist is selling. Traditional arguments by and large failed to motivate relativism over scepticism. And new epistemic relativism, of the semantic variety, to the extent that it's plausible – and as I suggested it that succumbs to several epistemological objections – was argued to be best understood as motivating not *epistemic relativism* at all, but rather (in Chapter 8) something akin to a Hazlett-style divorce between the epistemologist's concept of knowledge and the ordinary folk concept which (*a la* MacFarlane) is claimed to be assessment sensitive.

But then, and *thus*, the cliffhanger gives us no further suspense: we have no good reason to give the authority platitude constraining the conceptual limits of epistemic facthood the kind of (intra-perspectival) gloss we should give it *only if* we already have a reason to accept the epistemic relativist's story. But *now* – *viz.*, primary conclusion in hand – we lack any compelling reason to think that the kind of epistemic facts countenanced by the epistemic relativist, facts with *merely* intra-perspectival authority, fall *within* the conceptual limits of *epistemic facthood*; such facts after all fail to satisfy a *non-relative* reading of the authority platitude. And thus, we have a principled reason to reject that the epistemic relativist *is plausibly a metaepistemological realist* of any sort; the metaepistemological realist after all must countenance *epistemic facts*, facts that fall within conceptual limits as circumscribed by the content and authority platitudes. The kind of reading of the authority platitude on which the epistemic relativist could plausibly count as satisfying it is, in light of the primary conclusion, no longer on the table. And accordingly, the best way to think of epistemic relativism, then, is *not* as a 'variety of metaepistemological realism', but rather, as an anti-realist view, one which fails to countenance the existence of the kind of facts – epistemic facts – that one has to countenance to be any sort of metaepistemological realist. And now to get the affirmative answer to the incompatibility question, we just need to insert the final piece of puzzle: Chapter 1 argued at length that typical projects in mainstream

epistemology reveal, in practice, that they are tacitly committed to taking *for granted* a realist picture of the facts which are being disputed in first-order debates. This was the argument sequence that showed how revealed meta-epistemological commitments can be discovered by investigating the pragmatic presuppositions of first-order disagreements of typical first-order debates in mainstream epistemology – and an investigation of these pragmatic presuppositions revealed that the *shape* that first-order debates characteristically take makes sense only when a kind of realist picture of epistemic facts (what I called mainstream metaepistemological realism) is taken for granted. Not only then does, by the primary conclusion, it turn out that epistemic relativism (in neither its traditional nor new/semantics form) doesn't enjoy (from Chapters 2–8) the kind of philosophical motivation that would be needed to pose a threat to the revealed metaepistemological realism that lies in the background of typical first-order projects in mainstream epistemology, but even more: *that it doesn't* fuels a deeper point: that epistemic relativism is best understood as falling *outside* and as such *incompatible with* the very *kind* of metaepistemological realism that first-order projects are committed to taking for granted. Mainstream epistemology is not only *not* threatened by epistemic relativism, but given the nature of epistemic relativism, it's not clear that it *could be*, given what (e.g. as per Chapter 1) mainstream epistemology is committed to taking for granted.

### 9.3 An existential challenge?

One might wonder whether, in light of the affirmative answer I'm defending to the incompatibility question, arguments for epistemic relativism should be understood as simply *irrelevant* to mainstream epistemology. I want to close the book by suggesting why such a conclusion would be too quick. This is for two reasons – one that is I think straightforward, and probably obvious, and the other more complex. It will be the more complex and revisionary suggestion which I think is most interesting. But let's get the straightforward one out of the way first.

In short, it should be stressed that the incompatibility question gets an affirmative answer because the meta-epistemological commitments of typical first-order projects in mainstream epistemology *in fact* reveal themselves as ones which, in light of the shapes they take in practice, take for granted what I called mainstream metaepistemological realism. And this realism is, as we've seen, incompatible with epistemic relativism, which as I've argued is best understood as anti-realist along the

existence dimension: what the epistemic relativist countenances under the description of an epistemic fact falls outside of the conceptual limits of epistemic facthood.

But this incompatibility is importantly *contingent*. What the typical projects *are* in mainstream epistemology is itself contingent. I keep referring to Goldman versus Feldman, Moore versus the sceptic, as paradigmatic examples of debates in mainstream epistemology that are best explained as having the characteristic shapes they do in virtue of a background taking for granted that there epistemic facts (which respect the content and authority platitudes) to uncover. But I mention these examples only because in Chapter 1 I gave their cases special attention in defending this point. But as I qualified, I took these debates to be *representative* of the kind of shape taken by most typical first-order debates<sup>22</sup> and not interestingly different. Now: *perhaps*, over time, the debates that characterise mainstream epistemology could slowly *shift*. Perhaps, the second-order revealed meta-epistemological commitments could be ‘carried by the practice’,<sup>23</sup> like, as Wittgenstein (OC §248) puts it ‘foundation walls are carried by the whole house’. And one can’t rule out that, over a long period of time, perhaps the debates that will be characteristic of mainstream epistemology could slowly change – to continue the Wittgensteinian analogy, like sands on a riverbed,<sup>24</sup> so that the pragmatic presuppositions of these typical first-order debates in epistemology become very different. Given that we can’t rule out *a priori* that such a gradual shifting of focus might transpire, we should be willing to grant that, although epistemic relativism is *currently* incompatible with the kind of metaepistemological realism that lies in the background of typical first-order epistemological projects, it needn’t always be this way. And so we can define a limited kind of relevance for epistemic relativism for mainstream epistemology: if (for instance) arguments for epistemic relativism were to become stronger than they currently are, they could perhaps *begin to slowly shift* which debates are the ones of focus in mainstream epistemology, and in this process, new kinds of debates might take the place of the presently paradigmatic ones, where the new ones gradually start taking something like a relativist conception of epistemic facts for granted. I’m not at all convinced that anything like this will transpire any time soon, and in fact, just as a Wittgensteinian analogy can be used in support of this kind of story, so a Wittgensteinian analogy (at least on Pritchard’s (2012; 2015a) recent interpretation) can be equally wielded against it its possibility.<sup>25</sup> At any rate, this is the first suggestion for how – *despite* the conclusions drawn in this book – epistemic relativism might have a kind of potential

import for mainstream epistemology. And the answer is, put crudely, that mainstream epistemology would have to become a lot different than it is now.

But let's suppose things stay by and large as they have been. In what follows – and this is the more complex point – I want to sketch a revisionary account of how (in light of the conclusions drawn) we should view the most *significant* import that arguments for epistemic relativism should be understood as having for projects in mainstream epistemology. It's revisionary because it replaces an old (and I think unhelpful) way of thinking about this relationship – one I've argued against – with a different way of thinking about it. According to the received way of thinking, whatever 'threat' epistemic relativism poses to mainstream epistemology is best understood as:

- (i) *arising from traditional (e.g. non-semantic) arguments for epistemic relativism* (e.g. of the sort surveyed in Chapters 3–5), where these arguments are regarded as challenging
- (ii) the *material adequacy* of claims widely accepted in mainstream epistemology.

What I shall argue from here on out is that what the conclusions we've reached actually recommend is a very different, revisionary sort of picture. Here it is, in outline form: to the extent that epistemic relativism broadly construed stands to threaten mainstream epistemology *at all*, the threat is (in a sense I'll shortly explain) best understood as an *existential* rather than *material* one, and what does the relevant work motivating this existential threat is *new* (semantic), rather than *traditional* (non-semantic), epistemic relativism.

For the moment, let's put the existential point on the back burner. The first step in unpacking this revisionary picture is to say why – in light of what I've argued in Chapter 8 – new (semantic) epistemic relativism poses a threat *at all* to mainstream epistemology (in any respect), even if it's obvious enough why I think traditional forms don't. The line I defended in Chapter 8 was, after all, *not* suggestive of the thought that there is any residual threat – *to mainstream epistemology* – for new relativism to pose. The dilemma for new relativism posed at the end of that chapter was, to reiterate, that the *stronger* the case MacFarlane can make for the thesis that the ordinary folk concept 'knows' is assessment-sensitive, the *less reason* we have to think that the ordinary concept is what it is that *epistemologists* are interested in under the description of knowledge.<sup>26</sup>

I am not backtracking on this argument – I maintain that (i) MacFarlane’s ability to convince us that the ordinary concept of knowledge is assessment-sensitive; and (ii) reason to think that the ordinary concept is something other than what epistemologists are interested in under the description of knowledge, are *positively correlated*, and *that* they are positively correlated goes a long way to render mainstream epistemology insulated from the kind of relativist position that would materialise *were* the ordinary concept of knowledge to be vindicated as assessment-sensitive. Rather than to backtrack on this point, I want to (toward the end of motivating the revisionary line I want to put forward) point out that things are, in the wake of this dilemma posed at the end of Chapter 8 for the new relativist *vis-à-vis* mainstream epistemology, more complicated than they originally appeared to be.

To appreciate this point, it will be instructive to imagine a foe that we can call the *second-wave epistemic relativist*. The second-wave epistemic relativist’s ambition is to show how a *further sort* of epistemic relativism might be motivated on *precisely the kind of grounds* that I appealed to (e.g., the point about positive correlation, framing the dilemma) *against* the new (semantic) relativist, insofar as she stands to threaten anything we might call epistemic relativism in epistemology. We can imagine that this foe – the second-wave epistemic relativist – reflects on the dilemma raised against the MacFarlane-style epistemic relativist – and begins to speak, unsettlingly, as follows:

“Let’s grant that this dilemma is ‘valid’ – *viz.*, that (i) the compellingness of the case in favour of the claim that the ordinary concept of knowledge is assessment-sensitive – the case MacFarlane wants to make – and, (ii) reason to think that the ordinary concept ‘knows’ is something *other than what epistemologists are interested in* under the description of knowledge, are positively correlated. Now, *with this positive correlation in hand*, we might reasonably draw either one of two very different kinds of conclusions. Call the first conclusion *insulationism* and the second *second-wave epistemic relativism*. Insulationism is the conclusion gestured toward at the end of Chapter 8: that the positive correlation that characterises the Hazlett-style dilemma is evidence that mainstream epistemology is simply *insulated* from arguments aimed at establishing that the ordinary concept of knowledge is assessment-sensitive. Again, the pro-insulationist thinking was: the more those (MacFarlane-style) arguments are compelling in their intended aim, *the more reason we have to doubt that the ordinary concept of knowledge is what epistemologists are interested in*, and *a fortiori*, that

facts about the semantics of the ordinary concept of knowledge are facts that are relevant to the truth conditions for statements attributing the concept of knowledge that's of interest in epistemology. But once the positive correlation is conceded, we might just as well draw a very different conclusion, *second-wave epistemic relativism*. Second-wave epistemic relativism is the view one never *knows* something *simpliciter*; there is knowledge relative to the *standards apposite to the ordinary concept of knowledge* (which, if MacFarlane is right, will be assessment-sensitive) and knowledge relative to the epistemologist's standards, which are not assessment-sensitive. The Hazlett-style dilemma posed at the end of Chapter 8 favours second-wave epistemic relativism just as much as insulationism."

To be clear, I don't think the second-wave epistemic relativist's argument<sup>27</sup> is a good one, but responding to it will be instructive, I think, in showing where the residual – and as I'll suggest existential – threat to mainstream epistemology posed by new epistemic relativism *really* lies. The most obvious way to suggest that the positive correlation (which I've highlighted and which our envisioned second-wave relativist is trying to exploit) simply does *not* favour the kind of 'second-wave' epistemic relativism being floated here, is to point out that nothing like relativism follows from the suggestion that 'knows' might have multiple *distinct* meanings, and thus that knows is polysemous.<sup>28</sup> The appropriate response polysemy is *disambiguation* not relativisation.<sup>29</sup> And the absolutist takes disambiguation as par for the course, given the ubiquity of polysemy in natural languages.<sup>30</sup>

But once we dig our heels in with this reply – which I think we should, in response to the flimsy/tricksy argument of the envisioned would-be second-wave relativist – we are on our way to appreciating a much deeper worry, one which takes a bit of unpacking in order to present clearly. The starting point for the worry is an observation: there are in fact now many stories which, drawing inspiration from Edward Craig's (1990) influential genealogically motivated *practical explication* of the concept of knowledge, aim to vindicate with reference to practical interests and purposes the *need for the ordinary/folk concept of knowledge* as being one thing or another, and which then *theorise* about the nature of knowledge (and in particular about the truth-conditions for knowledge ascriptions) in light of conclusions reached. I want to now say a bit more carefully what is involved in these kinds of projects.

Firstly, some terminological clarification. As Kappel (2010, §1, 71–72) puts it, a *practical explication* of something X, answers the question: *what*

is the point of the X? – where X is some explanandum. A practical explication for explanandum E will have this kind of structure, a conjunction of a *need* claim and a *fulfilment* claim:

*Practical Explication for explanandum E:*

(Need claim): Given a set of facts F, and a set of aims or interests I, we have a certain need N;

(Fulfilment claim) E is what actually fulfils the need N.<sup>31</sup>

Craig's (1990) practical explication of the ordinary concept of knowledge identifies, in light of the relevant facts and our aims/interests, the *need* in question (N) as that of *tracking reliable informants*, and his position is that the concept of knowledge is what fulfils this need. In short, the primary point (alternatively: function) of the concept of knowledge is to identify good informants. More generally: practical explications thus answer the question *what is the point of something* by telling us, like Craig did, something important about the thing's function with reference to need-fulfilment.

A point of clarification is now needed. As I noted, practical explications have, at least in recent epistemology, been *used* to *theorise about knowledge* and in particular to inform semantics for knowledge-attributions. Robin McKenna (2013, 335–336) puts this point helpfully in perspective:

Plausibly, a successful analysis of the concept of 'knowledge' must match and predict our intuitions about a wide range of cases and explain why knowledge is valuable....Call this the 'value turn' in epistemology (Riggs 2008). Perhaps equally plausibly, *a successful analysis of knowledge must also fit with an account of the distinctive function or social role that the concept plays in our community [...]* Call this the 'functional turn' in epistemology.<sup>32</sup>

The relationship between a practical explication of the concept of knowledge and the (full-blown) *functional turn* in epistemology is this: participants in the functional turn in epistemology *appeal to practical explications* of the concept of knowledge, on the basis of which they identify a function, where that function is regarded as generating an *ex ante* constraint on an analysis of knowledge (or a semantics of knowledge attributions). Henderson (2009; 2011), McKenna (2013; 2014), Grimm (*forthcoming*), Pritchard (2012) and Hannon (2013; 2014; 2015) have for

instance defended views about the concept of knowledge (or knowledge ascriptions) that are more or less inspired by Craig's (1990) favoured account of the function of knowledge as identifying good informants.<sup>33</sup> By contrast, Kappel (2010), Kelp (2011) and Rysiew (2012) identify *closure of inquiry* as the relevant function and regard *this* rather than Craig's tracking-good-informants function as generative of an *ex ante* constraint for theorizing about knowledge and its truth-conditions. For Lawlor (2013) the relevant function is identified (*a la* Austin) as that of *providing assurance*, and for Beebe (2012), it's expressing epistemic approval/disapproval. So all else aside, there is a diversity of rival practical explication stories, ones which put their fingers on different functions.

Operating in the background of these function-first projects which appeal to some practical explication to identify a function *and then* use the function to inform semantics is an explicit (or implicit) commitment to thinking that something like the following is true, which is, as McKenna puts it, that:

on the best interpretation of talk of the function of a concept such as 'knowledge', the function of 'knowledge' is relevant to semantics.

The problem though, which Mikkel Gerken (2015) has recently highlighted, is that one can very easily take this idea too far. Gerken has cautioned against *overestimating* 'the extent to which this function should impact the semantics' (2015, 232) of knowledge ascriptions. Consider that, at one limit, we might assume that there is just *one* such function we're looking for (when giving a practical explication of the concept of knowledge) and that whatever that one function is, we should construct our semantics for knowledge attributions entirely around that function, such that the constitutive function of knowledge attributions is whatever the identified function is. This is more or less the move Lawlor has made in moving from an identification of the 'providing assurance' function and reasoning from this to the claim that the constitutive function of knowledge ascriptions is to provide assurance. Gerken (2015) is, I think, rightly critical of this move in his rationale for why moving as Lawlor does from 'a prominent functional role claim to a semantic conclusion is problematic', and in a way that I think suggests a more *general* lesson. As Gerken puts it, the

step from "common functional or communicative role" to semantic or substantive claims may not quite amount to a speech act fallacy (Lawlor: 42). None of the functional role theorists propose a full

semantics of 'knows' in terms of its communicative functions. Nevertheless, the step requires that the communicative roles be semantically constitutive. *But this is not plausible if there are classes of exemplary cases in which knowledge ascriptions serve a range of completely different conversational functions than the favored one.*

And Gerken thinks, plausibly, that knowledge ascriptions can perform a range of functions in exemplary cases. I think he's right about this. One obvious place to look is simply to the *other candidate practical explications*. There seems to be some truth, after all, to the tracking-reliable-informant approach, the closure-of-inquiry function, the expressing epistemic approval/disapproval approach as well as (and as Gerken grants) Lawlor's favoured providing assurance approach.

I mention Gerken's warning against moving seamlessly from practical explications to semantics in order to draw a sharp line between two very distinct *roles* a practical explication might be argued to play: *semantic* and *vindictory*. A practical explication of the concept of knowledge plays a *semantic* role just in case one attempts to inform one's semantics with reference to the practical explication. And, as Gerken points out, *some* ways of doing this are problematic.

But, there is an entirely different *reason* one might have for providing a practical explication for a concept, one which is entirely unrelated to any aims we have insofar as we might want to inform our semantics. This reason is what we can call *vindictory*. In short, rather than to answer the 'What is the point of X?' question *in the service* of informing our semantics of X-attributions (or, relatedly, in the service of offering a conceptual analysis of X), we might be in need of a practical explication of X for reasons broadly analogous to those a corporate downsizer has when asking various employees, 'What is the *point* of you here?' And, as it goes, if they can't give a satisfactory answer, they're regarded as expendable. In a sense, a practical explication when playing the semantic role (i.e. by the function-turners) is *forward-looking* whereas a practical explication, when serving a vindictory role, is *backward-looking*: the practical explication, when it plays the vindictory role, is used in the service of assessing whether the concept in question is actually needed because, looking backward, it can be compellingly defended as fulfilling some need.

We are now in a position to start putting several things together. Practical explications for the ordinary concept of knowledge yield an embarrassment of riches. This is one of the reasons that (*a la* Gerken) it's a bit dubious to try to make just *one* such alleged function the centrepiece

of one's semantics. In short, there are plenty of reasons to think that (and with reference to Kappel's need/fulfilment model of practical explication) there are multifarious needs such that the ordinary concept of knowledge fulfils them. Even more simply: the ordinary/folk concept of knowledge is one for which there is a point (or as it were *several* points) to having. In this respect, the ordinary concept of knowledge is *vindicated variously by the multiplicity of practical explications that have been offered for it*.

It goes without saying that if some concept (and norms associated with the concept) sustain a certain practice,<sup>34</sup> and that concept itself is simply not one for which we can give a very compelling practical explication – *viz.*, by identifying with reference to facts and purposes/interests, some need such that the concept can be demonstrated as fulfilling that need – then the practice is threatened in the following way: the concept the analysis of which sustains the practice is one for which we have no good answer to the 'What is the point of it?' question. And this is an *existential* threat to the practice, a threat whereby the rationale behind the existence of the practice is directly called in to question.

Suppose, for example, there was an unreflective practice centred around analysing some *grue*-like knowledge-related concept, *g-knowledge* – *viz.*, practitioners inherit this practice and unreflectively – that is without reflecting about whether there is anything about 'g-knowledge' which justifies the extensive attention to it – and attempt to work out necessary and sufficient conditions for g-knowledge, semantics for g-knowledge-attributions, etc. Now, suppose the g-knowledge-analysis practitioners learn a new fact, thanks to a kind of hired Craigian genealogists' depressing empty-handed *non*-story for 'g-knowledge': what the practitioners learn is that despite what they had assumed, there is really *no* compelling practical explication for the concept 'g-knowledge'; given the relevant facts and aims/interests, there is no obvious need (communicative/practical) such that 'g-knowledge' actually fulfils that need. The discovery of such a fact would leave the imagined practice centred around the analysis of g-knowledge existentially threatened in that it would be now be salient that it's not obvious what the point of the practice is, or why the practice should exist.

Now, putting the remaining pieces together: we can begin to appreciate why, *to the extent* new (semantic) relativism constitutes a kind of threat to mainstream epistemology at all, the threat is best understood as an *existential* one. The point can be spelled out as follows: *if* the new epistemic relativist (e.g. MacFarlane) offers the most compelling semantics for knowledge attributions on the table – more compelling than

competing views – then as I’ve suggested (e.g. Chapter 8), this is good reason to think that the ordinary concept of knowledge is something other than what epistemologists are interested in under the description of ‘knowledge’. And to the extent that we double-down on this kind of response – which as I suggested earlier in this section, I think we should<sup>35</sup> – then the following question becomes considerably tougher to answer than it would be otherwise: *what’s the point of the epistemologist’s (non-relative) concept of knowledge?*

And here’s why: on the familiar assumption that the *two concepts are the same* (or close enough to the same), the epistemologist can help herself unproblematically to the various practical explication stories that are already well established in the literature and which have been put forward in the service of vindicating the ordinary/folk concept of knowledge. That is, *on the assumption* that the ordinary concept of knowledge is what epistemologists are interested under the description of ‘knowledge’ – an assumption that is of course well entrenched given the extent to which the semantics of ordinary knowledge attributions are widely taken to be illuminative of the concept of knowledge epistemologists are interested in<sup>36</sup> – the epistemologist can reason that the concept of knowledge of interest to epistemologists is one there is a point to *because this concept helps us to \_\_\_*, where the \_\_\_ is filled in by the Craig-story, the closure of inquiry story, the assurance story, or whatever practical explication/function account one finds most appealing.

*However*, and this is the obvious catch: *if* the ordinary concept of knowledge is importantly different from the concept epistemologists are interested in under the description of ‘knowledge’ – something which in Chapter 8 I argued would plausibly be the case if MacFarlane’s argument succeeds – then epistemologists would *not* be entitled to help themselves to any of the established practical-explication vindication stories, stories which constitute practical explications of the *ordinary* concept of knowledge. Rather, as the line goes, epistemologists will need to find some other practical explication story for why the concept epistemologists are interested is needed/ useful. This is at any rate what I’m calling the ‘existential’ challenge that the new relativist can lay claim to, *even though* I’m answering the incompatibility question in the affirmative – and thus denying that epistemic relativism (of traditional and new/semantic varieties) is even compatible with the kind of metaepistemology taken for granted by most projects in first-order epistemology. In sum, I think that something like this kind of existential threat just described – one where the epistemologist loses any claim right to the genealogical stories that are appealed to in vindicating the need for the ordinary concept of

knowledge – stands as probably the most serious kind of threat epistemic relativism (of any variety) ultimately poses to mainstream epistemology, one that of course is pressing only to the extent that a truth-relativist semantics for knowledge attributions is difficult to put decisively to rest.

Of course, the principal *arguments* adduced against new epistemic relativism originally (in Chapter 7) were *epistemological* in character, and it's not comforting to consider that the force of such arguments is weakened if it turns out that the non-relative concept of knowledge that is of interest to epistemologists is not one which can be vindicated via practical explication. To be clear, the existential threat I've sketched does not *rebut* any particular vindicatory story for why a non-relative concept of knowledge fulfils some need that vindicates a practice sustained by analysing it. Rather, the challenge is best understood as *undercutting*: to the extent that the new relativist is effective in demonstrating that the ordinary concept of knowledge is assessment-sensitive – and to the extent that we reply by holding ground and insisting that this is further reason to *deny* that the ordinary concept is what's of interest to epistemologists – epistemologists have the usual-suspect practical explications for the concept of knowledge *swept out from underneath*, and must find new ones. What might these be? What, given the relevant facts and our aims/interests, is the relevant *need* the fulfilment of which is such that the epistemologist's (non-relative) concept latches on to it?

I want to end not by answering this on behalf of the mainstream epistemologist – my objective was simply to radically reposition how we understand the relativist's threat, to the extent that she poses one at all, to epistemology. Rather, I'll suggest where might be a promising place to look for to locate such a need, and that's to the very beginning of epistemology, as we know it – specifically, to Descartes. It's not implausible that the need for a sure foundation – regardless of whether this is easily attainable, and which guided Descartes' method of doubt, constitutes a need such that only the non-relative (epistemologists') concept of knowledge, and not the ordinary concept should MacFarlane be right that 'knows' in the ordinary sense is assessment-sensitive, suffices to fulfil that need. To the extent that this is on the right track, the existential challenge the relativist offers is none other than the challenge to reflect on the very reasons we begin doing *epistemology* in the first place, and whether these are good ones.

# Notes

## 1 Metaepistemology and Realism

1. Epistemology is (put very roughly) comprised of debates nowadays about (among other things) the following issues: the nature and scope of knowledge (including here the traditional projects of analysing knowledge and responding to the sceptic, broadly construed) as well as debates about justified belief, epistemic luck and Gettier problems, epistemic value, pragmatic encroachment, epistemic closure principles, intellectual virtues, rationality, basic knowledge, testimony, perception, memory, the role of intuitions, a priori knowledge/justification and epistemic norms. See Ichikawa and Steup (2014) and Pritchard (2004) for a helpful overview of some of these issues.
2. As Alston (1978, 276) puts it, in a way that will be contentious to some, are not 'susceptible to *objective truth values*' (my italics).
3. For some contemporary versions of this kind of view, see Huemer (2005), Enoch (2011), Schafer-Landau (2005) and Parfit (2013).
4. See, for example Goldman (1967; 1976) and Armstrong (1973).
5. The JTB analysis, in its simplest form, says that, necessarily, S knows that p just when (i) p is true, (ii) S believes that p, and (iii) S is justified in believing that p.
6. See, for example, Chrisman (2007), Field (1998) and Gibbard (1992). One outlying earlier defence of something akin to a non-cognitivist view of knowledge attributions was defended by Austin in the 1960s, a view challenged in Geach's (1965, 463) famous paper 'Assertion', the *locus classicus* of the Frege-Geach problem. Thanks to Matthew Chrisman for drawing my attention to this.
7. Though, note that other contemporary variants of the causal theory of knowledge include virtue reliabilist approaches to analysing knowledge (e.g. Greco 2010; Sosa 2009).
8. For the most rigorous defence of this position, see Conee and Feldman (2004).
9. Now, a well-known methodological objection that reliabilists have levelled against proponents of internalist JTB theories, such as evidentialism, is the following: an adequate theory of justification must explain justification in terms of some substantive conditions that are themselves not explained in terms of concept of justification or its cognates, 'or any concept closely aligned with justification, such as reasonability or rationality' Goldman (2011). In light of this requirement, if evidentialists cannot explain 'evidence' in non-justificational terms, then (given that reliabilists can by contrast meet this methodological condition), the reliabilist can claim that the evidentialist is not in the market to give a correct account of epistemic justification. But notice, crucially, that this objection – in short, a kind of circularity objection – is hardly one that (if correct) would imply that – regarding the factuality of claims about epistemic justification – the reliabilist and evidentialist have different metaepistemological positions.

10. Goldman sometimes uses 'justifiedness' rather than other cognates. But he uses the other cognates as well. See Goldman (1979).
11. Goldman himself would reject the claim that he is not giving an account of justification, construed evaluatively. He writes that 'The term "justified" I presume is an evaluative term, a term of appraisal... I want a set of substantive conditions that specify when a belief is justified' Goldman (1979, 1).
12. As a point of clarification, Alston's approach doesn't obviously predict that Goldman and Feldman would *recognise* any time their first-order disagreement committed them to a further second-order disagreement. Rather, the point is that if we follow Alston and understand Goldman and Feldman as (in virtue of their first-order disagreement) also disagreeing at the second-order in the way Alston suggests (along fact/value lines), then this is hard to reconcile with this with the way Goldman and Feldman actually debate – *viz.*, in a way that pays no attention to such dividing lines, and does not involve the claiming of the kinds of advantages (associated with these division lines) that would be expected.
13. For one thing, as Dancy (1982) notes, many contemporary intuitionists view the fact/value divide as misconceived, and so intuitionism should be in principle characterisable *without* reference to this claimed divide; further he remarks, as we've already seen Goldman (1979) himself suggest, that 'a naturalist might admit that "justified" is an evaluative term, but claim it is also factual.' See Dancy (1982, 397).
14. Dancy remarks that the most important criterion separating naturalists and intuitionists – the crucial dispute – 'is between the intuitionist's pluralism and the naturalist's monism; positions on the issue of definability are a consequence of positions on the monism/pluralism debate' (Dancy (1982, 397)). *Cf.* Maffie (1993) for a very different interpretation of metaepistemological naturalism, according to which 'intrinsically rational ends or principles' are eschewed, and more generally, questions about the nature and province of epistemology are viewed as scientific questions (1993, 1).
15. As Aidan McGlynn has pointed out to me, Dancy perhaps would no longer as he did in the early 1980s take this to be a commitment of intuitionism.
16. The reliabilist of course thinks that some beliefs – perceptually formed beliefs – will have the property of being justified when they have the property of being produced by perception, a reliable faculty. In this case, the 'natural base' is the belief's perceptual origin. But as a point of obviousness, this particular natural property is not claimed by any causal theorist to be shared in all cases of justified belief. Consider that, in his discussion of how basic beliefs might be immediately justified, Goldman (2008, 12) writes: 'Suppose introspective classification is a generally reliable process [...] Then any belief formed by introspection [...] will be justified'. Clearly, even theories that characterise epistemic justifiedness as a function of being caused in certain ways, do not strictly speaking equate the property of being justified with some particular, or single, natural property. In short, causal theories in epistemologies are simply not monists in the sense that Mill was a monist.
17. Dancy has in mind here, in particular, Chisholm (1977).
18. Even some avowed evidentialists rely indispensably on a causal condition. For a good recent example, consider the kind of 'explanationist evidentialism' defended in recent work by Kevin McCain (2013; 2014).

19. For a fuller articulation of the principle of veritistic value, see Goldman (1999). This position has also been referred to more recently as the thesis of epistemic value truth-monism.
20. *Ibid.*, p. 397. In a certain inclusive sense, all approaches to analysing knowledge – including virtue-based approaches – are ‘JTB’ theories – some just have very different conceptions of what the ‘J’ (or, indeed, further conditions beyond the ‘J’ and the ‘T’) are to look like.
21. Ichikawa and Steup (2014).
22. See, for instance, Goldman (1979; 2011). See also Pritchard (2008) for a methodological defence of this suggestion.
23. This is not to suggest that first-order disagreements can never connect with second-order disagreements; that’s obviously not the case. Rather, the idea, as I’ll develop in the next section, is that first-order disagreements take place against a shared background.
24. That is, a gloss of metaepistemology that parallels Sayre-McCord’s articulation of metaethics.
25. I say ‘in part’ because, metaepistemological commitments can also be articulated, as per the distinction drawn in §1.1.
26. In his book *Metaepistemology and Skepticism*, Richard Fumerton (1995) remarks that (for instance) ‘the arguments of the sceptic most often presuppose metaepistemological commitments’, and in this sense, Fumerton is (as I am doing) linking metaepistemological commitment with first-order positions (in this case, the debate between sceptic and anti-sceptic) by way of the notion of presupposition. However, that said, Fumerton is using ‘metaepistemology’ in a way differently than I do. On his view, metaepistemology is ‘first and foremost an investigation into the analysis of concepts central to epistemological thought’ Fumerton (1995, 3). I am not comfortable with this characterisation. ‘Knowledge’ is a concept central to epistemological thought, and the analysis of the concept of knowledge is associated with – in fact, it is the hallmark of – first-order epistemology. I thus prefer Sayre-McCord’s (and Cuneo’s) general view of metaepistemology to Fumerton’s – despite the wider sympathies I have with Fumerton’s compelling project in the book.
27. Suppose one were to insist that Feldman and Goldman, when arguing about justification and knowledge, are engaged in a merely ‘verbal’ dispute. Would such a suggestion hold water? It’s hard to see how. As David Chalmers has recently characterised (merely verbal) disputes: ‘A dispute over [sentence] *S* is (broadly) verbal when for some expression *T* in *S*, the parties disagree about the meaning of *T*, and the dispute over *S* arises wholly in virtue of this disagreement regarding *T*’ (Chalmers 2011, 522). On the presumption that the dispute is merely verbal, we should expect that Goldman and Feldman will (after arguing about these notions since the 1980s) have given serious attention to the possibility that they don’t actually disagree. This has not been the case. Consider here, further, Carrie Ichikawa Jenkins’s (2014) characterisation of merely verbal disputes, as follows: ‘Parties *A* and *B* are having a merely verbal dispute iff they are engaged in a sincere *prima facie* dispute *D*, but do not disagree over the subject matter(s) of *D*, and merely present the appearance of doing so owing to their divergent uses of some relevant portion of language’ (Jenkins 2014, 21). Goldman and Feldman are not engaged in a merely verbal dispute on Jenkins’s proposal any more

- than they would count as doing so on Chalmers'; they fail Jenkins's merely-verbal-dispute condition that they 'do not disagree over the subject matter of *D*'. I will be taking the disagreement between Feldman and Goldman to be a genuine one.
28. See here MacFarlane (2014, 121) for a statement of this kind of view, which he characterises as a standard view in contemporary philosophy, when formulated as both necessary and sufficient for genuine disagreement. Here we are only borrowing the necessity condition.
  29. There are also dynamic models of presupposition, though I won't be exploring those here. See Beaver and Geurt (2011).
  30. Beaver and Geurt (2011, §4.2).
  31. This is for two reasons. First, those aiming to reject one another's views don't always use conflicting sentences to do so. But, even if first-order disagreements could be understood in terms of patterns of conflicting sentences, the kind of presupposition that seems to connect with metaepistemological commitment is presupposition-cum-what-is-taken-for-granted, a notion that is not close to the simple notion of entailment, and much closer to the idea of pragmatic presupposition.
  32. See here Beaver and Geurt (2011, §4.2). For some seminal papers on pragmatic presupposition, see Stalnaker (1973; 1974; 2002). One important reason that pragmatic presupposition will be more helpful to us than semantic presupposition involves the famous 'projection problem' – the problem of accounting for how the presuppositions of complex expressions can be explained in terms of their sentential constituents.
  33. My italics. As Blome-Tillmann (2009, 253) characterises the view: *S* presupposes *p* in context *C* iff *S* is disposed to behave, in her use of language, as if she believed *p* to be common ground in *C*. Cf. Gerken (2012) for the view that presuppositions are non-attitudinal commitments.
  34. Note however the qualifier 'ordinarily'; as Rothschild (2007) remarks, 'A proposition need not actually be believed to count as part of the common ground, but the participants must at least pretend to believe it or must use it as a working assumption'. Stalnaker (2002) offers a helpful overview of some of the ways common ground and common belief can come apart.
  35. Here it is helpful to look to Grice (1991), whose thinking influenced Stalnaker on this point.
  36. The sceptic's conclusion is that there is no knowledge. Note, the epistemological sceptic is not denying that there is an external world. This, by contrast, is the position of the idealist. In rough template form, the sceptical argument to this conclusion is that: we know everyday propositions only if we know we aren't radically deceived, and further, that we don't know we're not being radically deceived.
  37. The closure principle actually admits of several versions. The core idea however is simple. If *S* knows that *p* and *S* competently deduces *q* from *p*, then *S* knows that *q*. See Pritchard (2005).
  38. The *locus classicus* of this reply to the sceptic is Moore (1925; 1939). For a diagnosis of 'neo-Moorean' responses to the sceptic, which take a premise like 'I know I have hands' as a starting point and then reason, via the closure principle, to the knowledge that sceptical hypotheses do not hold, see Pritchard (2002).

39. The 'begging the question' charge has been levelled by, among others, Lehrer (1971) and Stroud (1984).
40. Of course, there are other ways to frame the dispute. The relevant proposition could be 'that there is human knowledge'; Moore accepts this, the sceptic denies it. I chose the claim that Moore knows there is an external world because it is simple and avoids muddying the waters with the fact that – with regard to the proposition that there is an external world – the sceptic does not actually deny this. This is what the idealist denies. Note also, that (as Mikkel Gerken has pointed out to me) we might be inclined to think Moore and the sceptic disagree about the standards for a proof. This might be so, and Neta (2007) has offered some textual support from early Moore to support such a reading. However, Moore and the sceptic can disagree about multiple things. They at least disagree about whether Moore knows there's an external world. I'm focusing exclusively on this disagreement here.
41. In fact, the issue of just where Moore's proof goes wrong has been a topic of much contemporary interest. See Pryor (2004) and Wright (2007) for sophisticated diagnoses of where the proof goes wrong, diagnoses that correspond to a dividing line in the literature on perceptual warrant between dogmatism (Pryor) and conservatism (Wright). This discussion is revisited, albeit for different purposes, in Chapter 5.
42. If knowledge *weren't* taken by the sceptic to be factive, then by the sceptic's lights, knowing that p would be compatible with being deceived about p.
43. It's beyond what I can do here to give a full account of how to unpack 'relevantly'.
44. As Moore notes, Kant had famously thought that such a proof should be required. By contrast, Wittgenstein and Greco's reading of Thomas Reid are examples of thinkers that would deny that such a proof could be efficacious on this score. What's relevant to this debate is that Mooreans and their opponents behave, in their use of language, *as though* there is a fact of the matter, something to settle.
45. Moore expands here, suggesting how a proof might settle as well the matter of how many misprints are on a particular page of a book.
46. Passage quoted also in Vogel (2005), who disputes that Moore was actually trying to prove his conclusion, despite obvious appearances. Cf. Carter (2012) for a critique of Neta (2007).
47. It is a mistake to think the (Cartesian) sceptic is, in virtue of her scepticism, committed to denying that the issue of whether scepticism is true can be philosophically settled. The (non-Pyrrhonian, Cartesian) sceptic proceeds with the aim of settling the debate in favour of the sceptic. A good example of such a strategy is Unger (1971). The real opponent of the claim that the philosophical problem of scepticism could be settled is the philosophical quietist (Wittgenstein 1969). For an overview of this kind of position, see Virvidakis (2008). Somewhat peculiarly, Unger's very recent work (e.g. Unger (2012)), in contrast with his earlier scepticism, might well fall into the quietist camp.
48. Consider a potential (tricky) objection: take two metaphysicians arguing about whether there are such things as facts, and in doing so, making the same kind of dialectical moves that we find between Goldman and Feldman. It would seem odd to draw the conclusion that to make sense of their

- disagreement we need to posit a shared background commitment to facts about whether there are facts. I think though that that's exactly what we must do. The metaphysician reveals that she is taking for granted that there is at least one kind of fact: facts about whether there are facts. Compare: the Pyrrhonian is not committed to the existence of facts, but the Pyrrhonian does not argue (as this envisaged metaphysician does) in such a way that reveals any such presupposed commitments. Thanks to Aidan McGlynn for raising this case.
49. Sometimes such contradiction is followed by holding one's guns, other times by retraction; Lehrer (1971) for instance, used to think scepticism was true; now he thinks his former self was simply wrong about that – he retracts.
  50. Consider here Frege's remark that 'If something were true only for him who held it to be true, there would be no contradiction between the opinions of different people. So to be consistent, any person holding this view would have no right whatever to contradict the opposite view, he would have to espouse the principle: *non disputandum est*' (Frege (1979, 233), my emphasis). Cited also in MacFarlane (2014, 35).
  51. Western, as opposed to Chinese and East Asian astrologers, for instance, disagree about the principles of astrology; common ground to the dispute is that there is a causal connection between celestial relations and human fortune. This false bit of common ground must be in place for their arguments to proceed.
  52. See Gibbard (1992). See Cuneo (2007), Ch. 2, for discussion.
  53. The non-cognitivist rejects that moral claims are truth apt, and further rejects the relevant truth-makers, moral facts.
  54. One notable exception here is Williams (1996). Though Williams uses 'epistemological realism' to refer to a position that is much more specific than just a tacit commitment to epistemic facts. Here is not the place to engage with Williams' view. Another exception is Cuneo (2007), who models metaepistemological realism off of meta-ethical realism. We shall examine Cuneo shortly.
  55. As Devitt (1983, 77) remarks 'Realism... requires the objective independent existence of common-sense physical entities.'
  56. According to Sober (1982, 369), 'Realism is a declaration of *independence*' (my emphasis).
  57. Miller (2002, 2).
  58. Cf. Fine (2001) and Dreier (2004), who appeal to an 'explanation' axis in distinguishing realism. As far as I can tell, this axis isn't useful for distinguishing specifically *revealed* metaepistemological commitments, even if it would suffice for distinguishing articulated metaepistemological commitments.
  59. Though, note that (for instance, in the subject matter of ethics) the error theorist can allow that some of the characteristic objects exists (e.g. consider the statement 'Tim is good'); what they deny is that any moral properties are instantiated, and so they deny that there are any objects that bear the characteristic properties. The non-cognitivist by contrast, simply nips things at the bud a step earlier and denies that (e.g. in the moral case) moral discourse is even assertive.
  60. Also, a more complicated variety of anti-realism that is anti-realist in virtue of denying the independence claim is Dummett's (e.g. 1991) denial of semantic realism by denying *recognition transcendent* truth conditions.

61. As we'll see later, things are a bit more complicated.
62. This is an adaptation of Miller's (2012, §0) articulation of generic realism.
63. Craig (1998) captures this point in his remark that 'We are not speaking here of causal (in)dependence: the fact that there would be no houses if people had not had certain thoughts should not force us into antirealism about houses.'
64. As C.S. Jenkins (2005) remarks, 'it is not straightforward to say exactly which kinds of dependence are mundane.' We'll return to this matter, in this chapter and also in more detail in Chapter 9.
65. As Miller notes, a view can be 'more or less' realist.
66. The point can be made in terms of SPCG. Assume that the *denial* of the strong claim that epistemic properties are not reducible to natural properties – thus, assume that they are reducible in this way. The nearest worlds where we stipulate that Moore and the sceptic take this for granted are *not* worlds where their debate about whether K(W) is relevantly different.
67. Cf. Pettit (1991).
68. The most notable example of *metaepistemological non-cognitivism* is *epistemic expressivism* (e.g. Chrisman (2007); Gibbard (1992); Field (1998)). The epistemic expressivist reconceives of epistemic discourse, as something other than what the surface grammar would suggest. Whereas the epistemic error theorist insists that when we say something like 'Bob's belief is (epistemically) unjustified' we are making an *assertion* – and in doing so are saying something capable of being true or false, the epistemic error theorist denies this – and likens epistemic claims to expressions of (epistemic) approbation/disapprobation – expressions of attitude. Thus, while the epistemic expressivist's main interest is the nature of epistemic discourse (rather than the nature of epistemic facts and values themselves), an upshot of the view is that we can make sense of such discourse while simply doing away with any such facts.
69. My italics.
70. As Matthew Chrisman has pointed out to me, claims of epistemic modality (e.g. 'A might be F') and semantic necessity (e.g. 'all bachelors are unmarried') raise difficulties for a strict reading of Cuneo's line here. Regarding epistemic modals: if mightness is best understood as a qualification of an evidential perspective *rather than* a feature of reality, then it's hard to see how one's belief that 'A might be F' would attain positive epistemic status in virtue of representing reality aright. And likewise if semantic necessity is, as Chrisman puts it, an 'ossification of our policy of how to use words' *rather than* a feature of reality, then it's hard to see how my belief about the unmarriedness of bachelors attains positive epistemic status in virtue of representing reality aright. One response to the epistemic modals problem for the content platitude is to insist that claims of the form 'A might be F' count as representing reality aright (in the sense relevant to the content platitude) provided the qualifications are apt ones, in that they reflect an accurate take on one's own evidential position; and likewise, regarding semantic necessity, this kind of line would insist that my beliefs about propositions expressing semantic necessity claims count as representing reality aright – again, in the sense relevant to the content platitude – provided these belief accurately characterise some of the words we use, in light of how we use

them (or something like this). These envisioned replies reveal how we might opt for stricter or more flexible versions of the content platitude Cuneo is adverting to. I am open to the thought that Chrisman's examples pose serious problems for strict readings of the content platitude, while I maintain that the kind of responses I outlined indicate ways that more flexible readings of the content platitude could accommodate them. In what follows, the more substantive points I make, *vis-à-vis* Cuneo's approach to circumscribing epistemic facthood turn on the authority platitude rather than the content platitude. Thanks to Matthew Chrisman for pressing me on these points.

71. Cited also in Miller (2012, 14).
72. One notable line against epistemic error theory is that it is self-defeating. For a sustained discussion of this objection, see Streumer (2012).
73. See here Berkeley (1878) for a more general defence of 'empirical idealism.'
74. I am characterising this view in a way that is neutral about the complicated issue of how perspective features in a description of what makes a view relativist. This topic is introduced in Chapter 2 and pursued in some detail in Chapters 6, 7 and 8.
75. See Kusch (2010) for a more detailed discussion of why it makes sense to hyphenate such a property.
76. See Boghossian (2006a) for a discussion.

## 2 Global Relativism

1. Even more, philosophers are divided over whether, if this thesis were true, this would *involve* a difference in the way things *are*, never mind what shape such a difference might take. This deep divide, as we'll see, features both in commentaries on Plato's attempted refutation of Protagoras's brand of global relativism advanced in the *Theatetus*, as well as in more recent discussions of relativism by Boghossian (2006a) and Kölbel (2011).
2. For a recent and detailed exploration of this issue, see Spencer (2014).
3. For a more comprehensive discussion of these approaches to defining relativism, see Baghramian and Carter (2015 §1).
4. See Baghramian and Carter (2015 §1) and Baghramian (2004).
5. The denial of monism is a centrepiece of all views Boghossian (2006a) calls relativism; he attributes to any such view what he calls an essential *pluralist* clause, to the effect that competing viewpoints are regarded as equally valid.
6. See here also O'Grady (2002) for a defence of this kind of view.
7. For discussion on this point, see Baghramian and Carter (2015).
8. There is not profound disagreement between those who measure a sack of potatoes according to pounds, as opposed to kilograms. We could imagine of course a superficial disagreement, where one party says 'It weights 8.5' and the other party says, 'No, it weighs just 3'. Once both parties realise the first is referring to pounds, and the latter kilograms, though, no facts of the matter remain in contention, the dispute about what the weight of a sack of potatoes is *merely verbal* (e.g. see Chalmers (2011)). One might point out that it is possible for there to be a separate argument about which system of measuring is better, pounds or kilograms, which is not merely a verbal

dispute. But note that this is not a disagreement about *what* a sack of potatoes *weighs*.

9. Cited also in Spencer (2014, 2).
10. As Spencer (2014, 2) observes, a consequence of this approach is that relativism and *n*-ary-ness are incompatible; this is however a point Spencer goes on to reject in his criticism of this way of thinking about relativism.
11. As Kölbel (2011, 17) captures this idea, more generally: "The term "relativism" is indeed normally reserved for theses that are at least in some sense controversial or surprising or philosophically interesting..."
12. For the most comprehensive defence of relativism on semantic grounds – for a range of different domains of discourse – see MacFarlane (2014).
13. The point that relativism (whether local or global) should be regarded as fundamentally a thesis about truth, has been relied on by philosophers interested in relativism for very different reasons. For instance, it is a point Hales (1997) relies on in the service of rendering a consistent *logical* framework for global relativism. Likewise, as we shall see later in this chapter, Burnyeat (1976) appeals to the idea that relativism is a 'theory of truth' at a crucial juncture in his defence of a charitable reading of Plato's criticism of Protagoras's global relativist thesis in the *Theatetus*. For a sustained discussion on this suggestion that relativism is at root an alethic thesis, see Baghramian (2004).
14. Cf. some versions of expressivism and indexical relativism (see Kölbel (2011, 18)).
15. Kölbel (2011, 20). But as Kölbel (2011, 18) notes 'not everyone accepts that attributing a feature always yields a truth-evaluable claim or judgment', and so there is some scope to deny the direction of the entailment that moves from 'a is F' to 'it is true that a is F'.
16. Aesthetic relativism is local in this respect. Just because you are an aesthetic relativist (suppose, you think what counts as beautiful depends always on one's individual perspective or taste), it wouldn't follow that you can't maintain that claims about logic, epistemic standards and ethical principles are absolute and objective.
17. A qualification is needed here. Imagine: an especially liberal relativist might claim to be a *parameter pluralist* (where the domain of relativisation is not just one parameter, but multiple parameters  $P_1$  and  $P_2$  for instance, local cultural norms *and* regional cultural norms). This position, however, can generate contradictory verdicts in cases where  $P_1$  and  $P_2$  come apart. A parameter pluralist, to avoid such a contradiction, would require some kind of favouring rule to adjudicate which parameter takes priority. We can also envision *parameter disjunctivists*, who insist that the domain of relativisation is,  $P_1$  and  $P_2$ ; parameter disjunctivists would require similar kind of favouring rule to adjudicate *which* parameter is the operant one, and regardless of  $P_1$  and  $P_2$  would generate contradictory verdicts. Finally, we could envision *parameter relativists* who insist that which parameter is operant is *relative* to some further parameter. Though it is not clear how one could be a coherent parameter relativist 'all the way down'. For a related worry, see the Wittgensteinian-style objection raised later in this chapter.
18. Burnyeat (1976, 173).

19. As Meiland (1979) has suggested, it's really no accident that global relativism and the charge of (something like) self-refutation seem to go hand-in-hand. Imagine, for a moment, that you are arguing with a global relativist. You point out that some consequence, *C*, of the doctrine is implausible. As Meiland (1979, 52) writes: 'But of course the relativist will simply reply that those same consequences are acceptable to *him* (the relativist), thus giving a reply which is not only consistent with, *but in fact flows naturally from*, the relativist's own philosophy.' And given that the relativist can (and indeed must) reply in this way, to similar sorts of objections – *viz.*, where implausible consequences of the position are cited, premises are regarded as implausible or reasoning is claimed to be suspect – Meiland proposes that 'critics from Plato to the present have turned to the possibility of self-refutation as the *only viable form* of refutation against this position because it gives the promise of not being vulnerable to [the same kind of] reply'. Global relativism is in this way something like the 'king' in chess – stalemates aside, the only way you can legitimately try to *capture* this piece is through the damning act of checkmate – anything less, the king still stands. And just as there is no move in chess more damning to the opponent than a move that brings about checkmate, there is really no charge against a philosophical view that is more damning to that view than the charge that the doctrine refutes *itself* – it's no good to be wrong, but even worse to be wrong through the utter embarrassment of having one's *own words* (somehow) turned against one.
20. The *peritrope* actually appears in a sequence of three separate lines of reasoning against Protagoras. This passage constitutes the most famous and often cited element of the view. It is also the most controversial.
21. I say 'perhaps' because it's controversial whether the liar sentence has any content or merely the trappings of content. Thanks to Aidan McGlynn for flagging this point.
22. The Socrates example is noted by Kölbel (2011, 12).
23. There is probably a case that could be made to the effect that dialectical rules could include such things as Gricean norms. To the extent that this is right, we might be inclined to think that dialectical self-refutation is just a special case of a more general type of self-refutation which occurs in part because of some *norm* that is in place, where norms can include both dialectical norms assented to explicitly or tacitly, as well as *assertoric norms* which will in the default case be assented to tacitly.
24. In fact, this is very strange indeed, that Plato attributed to Protagoras the unrelativised premise. As Burnyeat (p. 174) notes, Vlastos (1956 p. *xiv*, n. 29) is puzzled as well on this point because, frankly, Plato should have known better. Vlastos writes that: 'Protagoras is very fussy about adding "for..." after "true" or "is" or "real"... While [Plato] puts the "for..." almost invariably while *reporting* or *describing* Protagoras' doctrine... he sometimes drops it in the course of *arguing* against Protagoras (e.g. in the [peritrope] argument in 171a), thereby inadvertently vitiating his own polemic'.
25. As Emilsson (1994, 136) puts it, 'Plato shows himself to be perfectly aware of the importance of the qualifiers elsewhere in the dialogue, it is hard to believe that their omission is a simply an error'. This is also Burnyeat's interpretation.

26. Note that Burnyeat reads Plato's 'seems' as substitutable with believes/judges. See p. 178, fn. 9. See here also Jordan (1971, 12).
27. As Burnyeat (1976, 178) notes, if 'all things are for each person exactly as they appear to him, no one can be mistaken about the reality that confronts him, of which he is the sole authoritative judge (measure)'.
28. Burnyeat (1976, 178) cites as textual support Plato's wind analogy at 152b, which is discussed later in this chapter.
29. Burnyeat (1976, 178). As Burnyeat puts it, 'Protagoras has to defend the equivalence of 'It seems to x that *p*' and 'It is true for x that *p*', not *merely an implication from the former to the latter*' (my italics).
30. As we'll see, the matter of what the proposition becomes relevant in the special case where *p* is the measure doctrine itself. But forget this point for now!
31. Note that this is *not* the much more controversial claim that Socrates does not believe that everything that he believes is true. For some of the issues that crop up in connection with this stronger claim, see Evnine (1999). What Socrates does not believe is that everything he believes is true *because* he believes it – *viz.*, that his very believing something makes something true.
32. This is a paraphrase of the formulation of the interim conclusion as described by Burnyeat (1976, 182). The only change is that I am using 'believes' for 'seems', which is a licit substitution by Burnyeat's own lights (p. 178, fn. 9). To be clear, (IC) follows (given NUT) simply from the contingent fact that Socrates in fact does not believe what NM implies for him (*viz.*, that anything he himself believes would thereby be true). See Burnyeat (1976, 179–80).
33. Meiland and Krausz (1982, 4) remark that: 'If all that were meant by saying a belief is true for Jones is that Jones holds that belief, then every belief that Jones holds would be true for Jones. But the relativist rejects this notion of relative truth; he or she takes the notion of relative truth more seriously than this ... just as our ordinary conception of truth allows a person to hold beliefs which are false, so too the notion of relative truth must allow an individual to hold beliefs which are false *for him or her*'. This is, of course, a *different* rationale for rejecting the reduction that Burnyeat thinks we should reject in characterizing Protagoras's view. Though, a wider point is relevant here. Meiland and Krausz's *rationale* for this point might well be one Burnyeat will be happy to accept provided we are thinking about the relativised truth predicate outside the context of subjective global relativism, where one's believing is a kind of truthmaker. The above passage is cited also in MacFarlane (2014, 40) in MacFarlane's setting up a philosophical problem a proponent of relative truth must address, which is that she must 'make clear why the relativized predicate she is explaining is a relativized *truth* predicate.'
34. Burnyeat (1976, 181). My italics.
35. My italics.
36. As Matthen (1985, 35) interprets the thesis, 'Private worlds contain public objects, but truths about these objects are private, and it is these private truths that constitute the private worlds.'
37. One could imagine a variant on the Heraclitian analogy, one on which one's *dispositional beliefs*, and not only one's occurrent beliefs, ground what holds in private worlds. It seems clear that the Heraclitian analogy takes for granted that the beliefs doing the truth-grounding work are occurrent. On such a

case, it's true for Socrates (at  $t$ ) that  $x$  is  $F$  if Socrates either judges  $x$  to be  $F$  at  $t$  or would be exposed to judge  $x$  as  $F$  at  $t$ . A puzzle that emerges though in cases where one judges contrary to one's disposition, a possibility that arises provided we allow that dispositions to judge don't require perfect reliability. In such an imagined case,  $S$  is disposed to judge  $x$  to be  $F$  but in fact judges  $x$  not  $F$ . In such a case, the claim that  $S$ 's dispositional beliefs ground private-world truth is at tension with the thought that  $S$ 's occurrent beliefs ground private-world truth. Thus, any variant on the Heraclitan analogy that attempts to allow dispositional beliefs to ground private-world truth will have to account for which beliefs, dispositional or occurrent, ground private world truth in cases where one's occurrent belief parts ways from what one is ordinarily disposed to believe. If not, private worlds can be both one way and not that way, at once.

38. As Burnyeat (1976, 181) puts it, 'And to borrow a timely remark that Socrates makes early on in the dialogue (152*b*), it is not likely that a clever man like Protagoras was merely waffling' as Burnyeat takes it he would have been if the equivalence were mere synonymy rather than 'an important discovery about our beliefs' (p. 181).
39. Note that this strategy, while for one thing flying in the face of the kind of claim that Protagoras submitted, as illustrated by the wind analogy, does not by any means 'guarantee' a realist interpretation of the PGR thesis.
40. One quickly loses a grip how a thesis can continue to be regarded as a *global relativist* thesis when envisioning the amputated version implied by option (i), on which statements like 'The wind is not cold for  $S$ ' (in a context where  $S$  feels the wind as warm) fail to count as following from the doctrine.
41. This is the argument strand at 170e7–170a1. The second argument in the sequence (171c5–7) is that, even if Protagoras believed his doctrine, since more people don't believe it than believe it, it would be (by the lights of Protagoras's doctrine) 'more false than true'. The peritrope is then meant as the nail in the coffin, establishing that the doctrine is (*a la* the first argument) not true for anyone, 'including the Sophist himself' (*ibid.*, 176).
42. My italics.
43. My italics.
44. In short, the strategy on which we show that Protagoras is dialectically stuck admitting IC\* and then, *via* the anti-realist unpacking of PGR, in conjunction with POT, show Protagoras to be stuck contradicting himself.
45. Nagel (2001, 15) echoes a similar kind of worry, which (like Hales) exploits what appears to be an unacceptable result that crops up once the global relativist thesis itself is (as it seems any *consistent* relativist must say) applied to itself.
46. Note that Hales himself is not *endorsing* this line of reasoning against the global relativist.
47. At the very least, the dialectical burden is then returned to the knock-down proponent, who must explain why, either, adverting to the translation principle is off limits *because* it involves adverting to absolute truths; or, alternatively, why the global relativist will end up adverting to some absolute truths which are objectionable by the lights of the global relativist's own theory *even if* the absolute truths the global relativist adverts to by way of the translation principle are not.

48. See Kupreeva (2014) for an alternative tension that is articulated in terms of an inconsistency between 'private' and 'public' rules for truth ascription.
49. Wright (2008b, 383). This anticipates an issue explored in depth in Chapter 6.
50. Though see Kusch (2010) for a different way to capture this idea.
51. See Chapter 7 for a more detailed discussion of this point.
52. The central target of *On Certainty* is however not the sceptic, but Moore (1939).
53. For helpful recent discussions of Wittgenstein's hinge epistemology, and its connection to sceptical as well as relativistic arguments, see Pritchard (2010; 2011).
54. Well, Kölbel rightly notes that there have to be *some* restrictions on what global relativism is supposed to apply to, to rule out that, say, contradictions, can be relatively true. Though PGR neatly gets around the need to make any significant concessions in the case of contradictions: the relative truths on PGR are the believed truths (given the biconditional), and believed truths will not, at least at any given time, be contradictions. This point of course is a controversial one.

### 3 The Pyrrhonian Argument for Epistemic Relativism

1. See Frede (1998a).
2. The locution 'the Pyrrhonian sceptic' is even somewhat misleading. For one thing, as Jonathan Barnes (1998, 60) notes, 'there can be no doubt that there was no single Pyrrhonian orthodoxy'. Here Barnes cites, for example, Galen's distinction between extreme and moderate sceptics *amongst* the Pyrrhonists of his own time. Galen drew this extreme/moderate distinction in terms of 'rustic' as opposed to 'urbane' Pyrrhonists, where the former withhold across the board (perhaps Pyrrho himself, according to his student Timon, was the only maximally rustic Pyrrhonist (Frede 1998a, 5) while the latter are drawn in to equipollence only in controversial areas, such as science and philosophy. The rustic/urbane distinction parallels another more contemporary distinction between, as Fine (2000, 206) puts it, the No Belief and Some Belief Views: do Pyrrhonian sceptics disavow all, or only some, beliefs?
3. Thanks to Modesto Gómez Alonso for helpful discussion related to this point.
4. See, for instance, Sankey (2010; 2011; 2012).
5. My italics.
6. This way of setting things up owes to John Greco (2013, 179).
7. This is the version of the argument as made explicit by Lammenranta (2008, 5); though note that Lammenranta claims that there are in fact three distinct versions of the Pyrrhonian sceptical argument, of which the regress formulation is but one. Lammenranta's line is actually that the regress formulation is not a credible sceptical threat. Rather, he insists the more potent version of the argument is what he calls the dialectical version of the argument.
8. Lammenranta's presentation, which I'm relying on, in order to be deductively valid, would need to make explicit a further premise or premise to make precise how it is that (1) and (2) would force one to accept (3).

9. Note that the Pyrrhonian sceptic would not purport to *believe* this conclusion; in fact, the Pyrrhonian Aenesidemus argued that Academic sceptics were dogmatists precisely because they claimed to know the sceptical conclusion. But even more carefully, it's unlikely (contrary to the distinction typically drawn between Pyrrhonian and Academic sceptics) that even most Academic sceptics would claim to *know* a conclusion such as (5). As Frede (1998b, 127) remarks: '[...] major ancient sceptics [both Academic and Pyrrhonian] were not concerned to establish or to defend any position, let alone the position that nothing is, or can be, known. In fact, they went out of their way to point out that, though they produced arguments for it, they did not actually take the position that nothing can be known. And they went on to criticize those who did claim that nothing can be known as being dogmatic as those philosophers who claimed that something can be known, as being pseudo-sceptics'.
10. The primary exponent of infinitism has been Peter Klein. See, for instance, Klein (2007; 1999; 2003) for some representative examples of Klein's case for infinitism as a better resolution to the regress problem than coherentism or foundationalism. Note, importantly, that Klein denies that infinitism entails scepticism – he regards infinitism as a non-sceptical resolution to the puzzle, though he admits that, for the infinitist, scepticism remains a serious possibility. For further discussion, see the introduction to Turri and Klein (2014, 1–19).
11. For an overview of the rationalist and empiricist commitments to foundationalism, see Turri and Klein (2014, 3–4).
12. See for example BonJour (1985). For an overview of the differences between holistic and linear coherentism, see Olsson (2014).
13. On this proposal a belief is justified in virtue of a chain that has the properties of being (i) never-ending *in the sense that* for each belief in the chain there is a reason for that belief also in the chain (Olsson 2014); but (ii) with a finite number of unique reasons in the chain. Together (i) and (ii) imply that at least one reason in the chain is repeated.
14. Lammenranta (2008) considers this possibility briefly. See also Williams (1991; 2007) for an appeal to the Wittgensteinian line under the description of Wittgensteinian contextualism. Though note that Williams' target – the metaepistemological realist – is an entirely different (and much more specific) target than the view I am calling metaepistemological realism in Chapter 1 and again previously in this chapter.
15. See, for instance, Aikin (2011; 2005). *Cf.* Aikin (2014), for an attempt to show how one might embrace the infinitist premise (3) alongside a *non-sceptical* conclusion.
16. *Cf.* however, D. Greco (forthcoming) for a recent contextualist alternative to traditional construals of foundationalism and coherentism.
17. Sankey (2011, 564.) By 'proposes a criterion' I take it Sankey means offers conditions that are satisfied just when knowledge is present.
18. While Sankey, in each of his papers, situates the regress within the language of the Problem of the Criterion (a problem in epistemological methodology typically associated with Roderick Chisholm (1973)), it's not clear (to me at least) why details specific to the POC should be regarded as distinctively important for the purposes of running an instance of the regress argument

where the  $p$  in question happens to be a claim that features importantly in Chisholm's problem. One way to make this point is simply point out that there is a shortcut to the very same instance of Agrippa's Trilemma that avoids entirely the philosophical import of the Problem of the Criterion. The shortcut arises when one simply says that criterion 'C' is a criterion for knowledge. The sceptic then asks for justification. But this request for justification is *enough* on its own to get Agrippa's Trilemma up and running. There is no need to point out, *in order to generate this instance of the Pyrrhonian regress argument* that additionally, as Chisholm has famously noted, the particular endeavour of specifying a criteria for knowledge brings with it a further 'catch-22' (a catch-22 that, actually, is itself just an instance of a more general argument pattern often deployed by Socrates). In short, as I see it, the 'catch-22' typically associated with the Problem of the Criterion is a third wheel for the purposes of getting to the point Sankey ultimately gets to – *viz.*, where we find ourselves run through Agrippa's Trilemma and reach its sceptical conclusion, and then looking back for a way out.

19. Of course, this is simply not taking the externalist seriously. I return to this point in some detail.
20. This is, note, a restricted way to frame epistemic relativism in that the object of relativisation (see the co-variance definition, Chapter 1) is limited to only one kind of epistemic fact.
21. See Sankey (2010, 3). Consider Sankey's definition in a wider context. Generally speaking, a norm is, as Srinivasan (2014) puts it, 'a universal generalisation about how one ought to conduct one's practical or doxastic affairs, involving a normative state N and a triggering condition C, of the form "N if/and/only if C".' Sankey's claim that an epistemic norm is 'a criterion or rule that may be employed to justify a belief' can be understood in terms of this more general template; a rule such as 'if perceptual evidence favors  $p$ , believe  $p$ ' can be employed to justify a belief, and fits the general template "N if/and/only if C".'
22. Compare, though the suggestion that we abandon talk of sets of beliefs for talk of perspectives, where perspectives needn't be thought of as sets of beliefs from which entailment relations can be generated (e.g. Hales 2006).
23. See Sankey (2011 §3, esp. pp. 564–566).
24. I am closely paraphrasing, but also reconstructing for clarity, the argument thread found in Sankey (2012, 187).
25. Sankey (2012, 187).
26. To be clear, this is Sankey reasoning on behalf of the relativist.
27. See here also McCain (2014).
28. As noted earlier, Sankey regards the Pyrrhonian argument outlined as the strongest way to capture the epistemic relativist's argument strand, and his own naturalistic anti-relativist response as best understood as an *overriding* anti-relativist strategy, in that the strategy recognizes the Pyrrhonian argument for epistemic relativism as a *legitimate* one and then attempts to overcome it. The strategy I'll be pursuing is, unlike Sankey's, *undercutting*. I will be suggesting that the Pyrrhonian argument doesn't motivate epistemic relativism in the first place. Cf. Pritchard (2015) for a related discussion *vis-à-vis* undercutting and overriding anti-sceptical strategies.
29. I am following Turri and Klein (2014, 6–7) in this distinction.

30. Cited also in Turri and Klein (2014, 6).
31. My italics.
32. Epistemic circularity and its role in arguments for epistemic relativism will be the centrepiece of Chapter 5.
33. Compare: objecting to panpsychism in the philosophy of mind shouldn't proceed by simply pointing out that panpsychism is committed to something like the omnipresence of the mental; this is exactly what the panpsychist wants to say.
34. This is a reconstruction of the objection as presented by Turri and Klein (2014, 12).
35. Turri & Klein (2014, 13) attribute this observation to S. Wright (2013 §3).
36. At least, so long as having an infinite number of reasons *available* does not entail possessing an infinite mind. One might be inclined to think that I can't have an infinite number of reasons available if I have a finite mind. After all, as this line would go, I cannot avail myself to this many reasons. However, I think that this retort misses Turri & Klein's point. *Availing* oneself to an infinite number of reasons is tantamount to producing them. The notion of having reasons available to one was introduced in contrast with actually producing them. (Compare: I can have available an infinite number of expressions in virtue of my competence in English even though I cannot make an infinite number of expressions, given my finite mind and life). Thanks to an audience member at the University of Vienna for raising this point.
37. See Sankey (2010, 5).
38. Sankey (2012, 187), my italics.
39. Sankey (2011, 566), my italics.
40. Or, more carefully: he is taking it for granted that the relativist, in putting this argument forward, is convinced that Premise (1) in the Pyrrhonian Regress argument is obviously true.
41. According to Seidel (2013, 137), the dilemma is this: 'the epistemic relativist using the Pyrrhonian strategy [...] either [...] goes a long way with the sceptic but, in turn, cannot say why we are epistemically justified *at all*, or she tries to avoid scepticism by relying on an epistemic meta-criterion that does not seem to be justified *relative* to her own community'. Sankey (2013) has replied to Seidel's argument. In effect, Sankey thinks Seidel is mistaken to suggest the relativist, having travelled as she does with the sceptic, 'cannot say why we are epistemically justified at all'. Important to Sankey's reply is the distinction between weak and strong justification. As Sankey (2013, 142) puts it: 'Let us define weak justification as justification of a belief on the basis of a given epistemic norm, whether or not the epistemic norm is itself justified. Let us define strong justification as justification of a belief on the basis of a given epistemic norm, where the epistemic norm is itself justified. In light of the Pyrrhonian problem of the criterion, we may conclude that there is no such thing as strong justification. This is because it is impossible to provide an epistemic norm with justification, so it is not possible to justify a belief on the basis of a justified epistemic norm. However, we might still say that it is possible to weakly justify a belief, since it is possible to appeal to a norm to justify a belief. If the belief complies with the norm, then the belief is weakly justified on the basis of that norm.' In short, while I think this is a creative and interesting reply on Sankey's behalf, I don't think it ultimately rebuts

Seidel's point effectively. The problem, as I see it, is that Sankey's adverting to what he calls weak justification, as a form of justification, conflates a distinction, owing originally to Peter Geach, between *attributive* and *predicative* uses of the modifier 'justified'. Once this distinction is clear, Sankey's appealing to weak justification betrays itself as a kind of equivocation when appreciated in the context of replying to Seidel's objection, which is *itself* framed in terms of a predicative rather than attributive use of 'justified'. As Ridge (2013, 188) summarises the distinction: "sometimes a locution of the form 'is an F G' entails both 'is F' and 'is a G', whereas in other cases this entailment does not hold. Geach calls uses in which this entailment does hold 'predicative' and uses in which the entailment does not hold 'attributive'. The locution 'is a dead fly' is predicative, since it entails both 'is dead' and 'is a fly'. By contrast, the locution 'is a big fly' is typically attributive, as it does not entail 'is big'. Something can be a big fly without being big, full stop." Just as a 'big fly' does not entail 'big' as an attributive rather than predicative modifier, the notion of a 'justified belief' *as used in what Sankey is calling the weak sense* does not entail that the belief is epistemically justified (in the sense of attaining the status of epistemic justification), only that someone has *attempted to justify* the belief by showing how the belief complies with some rule. So what Sankey is calling weakly justified beliefs uses 'justified' in merely the attributive sense, where plausibly Seidel's challenge at (2013, 137) is one that, in noting that Sankey's relativist travels so far that she (as Seidel puts it) 'cannot say why we are epistemically justified *at all*' is best read as using justified in the *predicative* sense. In this respect, Sankey's preserved 'weak' justification in reply to Seidel looks like a form of equivocation. Though see also my connected point in §3.3, in the discussion of the Poison Oracle case.

42. If epistemic relativism (on Sankey's definition) is true, and so if there are no epistemic norms (criteria or rules used to justify beliefs) over and above the variable epistemic norms operative in different (local) cultural settings or contexts, then since norms just *are* rules or criteria used to justify beliefs, beliefs (if they are justified at all) will be justified only with reference to locally adopted criteria or rules. And so statement (ii) is a corollary of epistemic relativism, as Sankey defines it. Thanks to Markus Seidel for requesting clarification on this point.
43. It is worthwhile to consider an alternative line of argument, suggested to me by Martin Kusch, by which the epistemic relativist could draw from the Pyrrhonian problematic but in a way that avoids the brunt of my challenge to the strand of argument Sankey attributes to the relativist. The key idea is that *relativistic construals* of foundationalism, infinitism, coherentism, etc., which offer a comparatively more lax way of thinking about the structure of justified beliefs than do absolutist construals of these approaches, open up ways to avoid scepticism not available to the absolutist. They do this by allowing justification to be attained more easily than it is on absolutist construals of these approaches. The primary problem I have with this suggestion is that it begs the question against the sceptic. This reimagined version of the argument fails to decisively give us reason to embrace relativism rather than scepticism. In that respect, it is not ultimately better off at Stage 2 of the argument, even if it has additional resources to avoid objections I raised to Sankey's relativist's way of defending Stage 1.

44. Though, see the Wittgensteinian point at the end of Chapter 2 for a reason to be sceptical that even on the Protagorean version of global relativism *everything* can be relativised in such a way that the ‘no-mistakes’ direction of the global relativist’s biconditional holds.
45. See, for an example of some illuminating discussions on this point, Barnes (1998), Frede (1998a), Frede (1998b), Fine (2000).
46. This is Lammenranta’s terminology. I take the locution ‘working in unison’ to refer to the overarching Agrippan sceptical strategy within which each of the modes has some role to play.
47. The obligation is not one the sceptic has, but the dogmatist, at whom the sceptic directs the argument.
48. The suggestion here is just that some cases will have these properties.
49. Cited also in Lammenranta (2008, 12–13).

#### 4 Dialogic Arguments for Epistemic Relativism

1. This is a somewhat simplified picture. Technically ‘relativity’ is among the five (not four) modes Sextus adverts to. As Lammenranta suggests, it is supposed to work with the mode of disagreement in setting up the challenge which cannot be completed. But Lammenranta admits that what is meant by ‘relativity’ is not clear in Sextus’s writings, nor is its precise function clear. This ambiguity of the role of the mode of relativity is echoed in a discussion of the challenge by Katja Vogt (2014), who writes: ‘Skeptical examination often begins with the *Mode of Disagreement*: different answers to a given question are surveyed, and the conflict between them is observed. The interpretation...hangs, for the most part, on the question of whether anepikriton should be translated as ‘undecided’ or ‘undecidable’ (Barnes 1990). It would be dogmatic to claim that matters are undecidable. The Pyrrhonist must prefer the idea that, up to now, matters have not been decided. This leads to the question of whether something can be found that would decide matters, and thus to the application of further modes. Scholars have observed that... the *Mode of Relativity*, does not really fit into the Five Modes’ (my italics).
2. Lammenranta (2008, 15–16).
3. From a 12 April, 1615 letter to Paolo Foscarini.
4. Even more, as Hales sees it, ‘the most promising candidate on the relativist ticket is that of disputes involving irreconcilable differences’, more promising he thinks, than are disagreements about *taste* (2014, 73–77), which MacFarlane (e.g. 2014) among others have appealed to in motivating relativism.
5. See Baghramian and Carter (2015) for an overview of what is *involved* in a given relativist position. See also Ch. 2, §2.
6. This is being presented here as at least a necessary condition for anything properly called an epistemic relativist resolution; though I’m not suggesting here it is sufficient. Later in the chapter, I return to this point.
7. See however Siegel (2004) for a more detailed presentation.
8. Note that Siegel himself in rejects this argument strategy as compelling, for reasons we’ll engage with later in the chapter.
9. My italics.

10. What about a meta-norm such as 'Believe what the evidence supports'? Galileo and Bellarmine might plausibly both *claim* to be following what the evidence supports. Superficially, it appears that they share this metanorm in common. However, while there is a sense in which they 'share' this metanorm, in so far as they might both be inclined to sincerely use the sentence 'Believe what the evidence supports' sincerely, it's important to note that Galileo and Bellarmine will unpack this idea very differently. They after all have a very different idea of what it is for something to be supported by evidence. Thanks to Robin McKenna for raising this issue.
11. Consider, a meta-norm that one would accept only if one already accepted Scripture may well discriminate in favour of Bellarmine's position just as a meta-norm that one accepts on the basis of scientific evidence would discriminate in favour of Galileo's position, though both instances fail to be appropriately neutral, despite being very discriminating.
12. Arguments where no Archimedean meta-norm is available for either side to appeal to will be arguments where, as Lynch (2010), puts it there will be ones where 'the spade has turned', such that the process of actively justifying in the epistemic sense loses its point (p. 273). Lynch is understanding the properties of deep disagreements, in this sense, as ones that have certain features which align very closely with the features a disagreement has when there is no available Archimedean meta-norm. See Lynch (2010, 265) for Lynch's preferred characterisation.
13. See Bengson and Moffett (2011b) for a more detailed discussion of how to distinguish these positions.
14. For some recent discussions of how these first-order positions tend to line up with these second-order views about the kinds of evidence most relevant to establishing intellectualism/anti-intellectualism, see Glick (2011) and Brown (2013).
15. Toribio, (2008, 43–44); quoted in Stanley (2011, 171).
16. Second italics are mine.
17. To be fair, Stanley has, since his 2011 book in which he engages with Toribio's argument, attempted, in a co-authored paper with neuroscientist John Krakauer (e.g. Stanley and Krakauer (2013)), to suggest that some cases in cognitive science positively support the intellectualist position that he reaches independently on semantic grounds.
18. More recently, *epistemological* evidence has gained relevance in this dispute. See for example Poston (2009), Carter and Pritchard (2015) and Cath (2015).
19. For an interesting clash between linguistic and phenomenological arguments as they bear on know-how, see Stanley (2011 Chs. 1 and 7).
20. See also Dreyfus (2007). For another recent example where phenomenological evidence – in this case from Merleau-Ponty – has been brought to bear on the know-how debate, see Berendzen (2014).
21. My italics.
22. See, for instance, Sankey (2010) and Siegel (2011).
23. See, for seminal discussions of this line of argument, Harman (1999); Harman (2000); Doris (2002). See also Alfano (2013) for a more recent presentation of the argument.
24. See especially Alfano (2012) and Alfano (2014).

25. See Carter and Pritchard (2015).
26. One area where the availability heuristic has been especially well-documented is in the literature on risk perception. The more exposure one has to some perceived risk (e.g. ebola) the higher one is likely to estimate the probability of the risk (e.g. Slovic (1987)).
27. To think otherwise is to make the notorious *conjunctive fallacy*, the fallacy which violates the fact that for an event A, the probability of A and B cannot not be greater than the probability of A.
28. See for example, Moss-Racusin et al. (2012) and Bertrand and Mullainathan (2003) for some example studies.
29. See Saul (2013, 245–246).
30. Two points are worth noting here. Firstly, is drawing attention to cognitive biases, in a certain sense, self-undermining for one who is employing empirical methods? After all, doesn't establishing empirical conclusions about the prevalence of cognitive biases at the same time indicate that the very empirical results (concluding that we are subject to various biases) are *themselves* the product of some biases and should not be trusted? The results however are not ones which conclude that the scientific methods (with extensive controls in place) for discovering biases are themselves biased – rather, they conclude that ordinary belief-forming mechanisms are. The second point, however, is that we can imagine (for example) Diane suggesting that, at least from her perspective, she is not biased – because the revealed word of god is free from bias. And more generally, we can suppose that John and Lise will very likely *deny* that biases have had anything to do with their reaching the irreconcilable position that they reached. Again, though, this is just what we'd expect. The empirical evidence suggests that many such biases, and especially situational factors that influence judgment, fly under the radar. And so the fact of an agent's believing she is bias-free counts for naught. Ultimately, the claim that, plausibly, when we reach irreconcilable positions in actual dialogues, it's at least in some part due to cognitive bias, is a claim that is supported by evidence about the pervasiveness of biases, and so resisting this claim would involve going against the evidence. Thanks to Steven Hales for discussion on these points.
31. Is it possible that, in at least some circumstances, what epistemic principles we embrace are partially determined by our cognitive biases? It seems plausible that this would be so. For example, one kind of cognitive bias, the 'bandwagon effect' a version of which is the 'informational cascade' (e.g. Bikchandani et al. 1992) occurs when individuals follow the crowd without regard to their own information. Now consider this comparison. In a 2006 survey by Eurispes, 87.8% of Italians identify as Roman Catholic. ([http://www.corriere.it/Primo\\_Piano/Cronache/2006/01\\_Gennaio/17/cattolici.shtml](http://www.corriere.it/Primo_Piano/Cronache/2006/01_Gennaio/17/cattolici.shtml)) By contrast, in the 2011 census, 0.41% of Estonians identified as Roman Catholic. <http://pub.stat.ee/px-web.2001/Dialog/varval.asp?ma=PC0454&lang=1>. This is a striking difference. As we saw in the case of Bellarmine, epistemic principles can potentially be engrafted in to the religious doctrines one accepts. Putting this together: the bandwagon effect/informational cascade might well explain why, for instance, some Italians are more inclined to certain epistemic principles than are Estonians, and vice versa. This is, however, is exactly as we should expect, and is compatible with

the particular point I am relying on about the relationship between cognitive biases and epistemic principles in the argument I'm advancing here, which is that they are not the *same* thing. Which epistemic principles one subscribes to is a matter, at least, of what beliefs one has about, for instance, what kind of things epistemically justify other things. (I think the way Boghossian articulates what an epistemic principle is plausible.) And the point I'm relying on is that what principles one accepts is a matter that is different from what cognitive biases one succumbs to. This should be clear in the John and Lise case, one in which two individuals who accept the same epistemic principles can succumb to a plethora of very different cognitive biases which lead them to an irreconcilable position. That cognitive biases can play a causal role in what epistemic principles we embrace does not undermine the key point that these are different kinds of things.

32. A clarification: we must distinguish the efficacy of *ordinary* reliance on possible cases, in the service of challenging a universal generalisation (the efficacy of which is not in question here) from the kind of possibilist strategy I've suggested a proponent of a dialogic argument for epistemic relativism must retreat to: one where it is merely possible disagreements between possible agents that are relied on to *motivate* epistemic relativism. Whereas the role of possible cases in the counterexample role *vis-a-vis* a universal generalisation is one of straightforward entailment, the same is not so in the motivating role. Accordingly, the worries raised toward a possibilist strategy such as the one the dialogist proponent of epistemic relativism must retreat to are *not* regarded as worries that would apply *mutatis mutandis* to the more general strategy of relying in possible cases in the counterexample role.
33. And, *possibly*, this property could be entirely a function of possible agents' subscription to different moral systems, systems which make injunctions about what is permitted in dialogue.
34. For a representative sample of recent work on the topic, see Feldman and Warfield (2010).
35. See, for instance, Christensen (2009), Elga (2007) and Feldman (2007).

## 5 Incommensurability, Circularity and Epistemic Relativism

1. Lynch (2010) refers to situations, so described, as 'deep' epistemic disagreements.
2. See Oberheim and Hoyningen-Huene (2013) for an overview.
3. See Sankey (1999, 2) for further discussion on this point.
4. It will be helpful to note one 'larger-picture' issue to do with incommensurability strategies, one that connects with my criticism of Sankey's attempt to deploy a Pyrrhonian Argument, *via* Agrippa's Trilemma, in the service of motivating epistemic relativism – a criticism which was the central focus of Chapter 3. An obvious commonality between Sankey's version of the Pyrrhonian strategy, and incommensurability arguments (as I've just described the contours of their structure) is recourse to epistemic *circularity*. Epistemic circularity is a core *element* of Sankey's Pyrrhonian strategy, and it's a strategy I argued didn't work. Do the arguments against Sankey's

redeployment of the Pyrrhonian argument apply, *mutatis mutandis*, for other strategies which attempt to reason from epistemic circularity to epistemic relativism? In short, the answer is no. Here is a summary of the role circularity played in the strategy Sankey envisions the relativist using (under the description of the Pyrrhonian argument for epistemic relativism). Recall that the wider Pyrrhonian argument (at least, as appealed to by Sankey's relativist) came in two stages: in 'stage one', the Agrippan modes of hypothesis, circularity and infinite regress are meant to show that, once one attempts to provide epistemic justification for an epistemic norm, the process of justification cannot be satisfactorily completed. The combined appeal to the three Agrippan modes was supposed to generate the result that, therefore, all epistemic norms are equally *unjustified*. This is the intermediate conclusion. The second stage of the argument, as I interpreted it, was to move from the alleged equal justification of all epistemic norms to the conclusion that epistemic relativism is true. So circularity features, alongside hypothesis and infinite regress, specifically in stage one of the argument. My central line of criticism against the first stage of the argument Sankey attributes to the Pyrrhonian-motivated relativist was, in summary form, that Sankey has the relativist take it as a given that an epistemic norm is satisfactorily justified only if (i) a justification is *provided* for the epistemic norm (a requirement that rules out, *ex ante*, various externalist and foundationalist strategies), and further that that justification provided is neither circular nor 'infinite' (which excludes, without argument, any attempt to defend non-viciously circular justification for epistemic norms (e.g. Boghossian (2001) and Sosa (2011)), or for that matter, epistemological infinitism (e.g. Klein (2007))). My criticism of Sankey's view of the relativist's attempt to appeal to circularity (within stage one of the argument just summarised) was thus, in a nutshell, a 'lack of defence' objection. The suggestion that adverting to an epistemic norm as a part of one's justification for that epistemic norm 'is to reason in a circle' *might well* be true, and further, it might well be 'bad', but a further explanation would be needed. This lack of defence objection of course did not exclude the possibility that there *is some circularity-based line of argument that stands to motivate epistemic relativism*. And it might well be that the best way to get such an argument up and running will be to begin with cases where, as in the original position, we encounter epistemic frameworks very different from our own.

5. As Kahane (2011) frames such arguments, they combine: (i) a causal premise to the effect that *S*'s belief that *p* is explained by *X*, with (ii) an epistemic premise to the effect that *X* is an 'off-track' process, to the conclusion that (iii) *S*'s belief that *p* is unjustified.
6. Of course, as Srinivasan (2014, 4) notes, 'not every revelation of genealogical contingency undermines judgment. My judgment that Paris is the capital of France is contingent on the fact that I exist at all, that I possess the concepts Paris and France, and that I have been taught that the capital of France is Paris. And yet none of these revelations of genealogical contingency seem to undermine my claim to know that Paris is the capital of France. If they do, we have entered a realm of wholesale scepticism, in which none of my judgments are secure.'
7. Cited also in Srinivasan (2014, 4).

8. According to Boghossian (2006, 101), provided I'm entitled to employ the principles of my own system without first having to justify these principles (what Boghossian calls 'blind entitlement' (p. 99)), the mere *encountering* of another epistemic system needs to have a certain amount of credibility in order to bring us to 'legitimately...doubt the correctness' (p. 100) of the system we're already using. As Boghossian sees it, the alternate epistemic system must 'have to be a real-life epistemic system, with a proven track record, not just some theoretical possibility. Its *actual* achievements would have to be *impressive* enough to make us legitimately doubt the correctness of our own system' (p. 101). While I think this requirement is too strong, and I make this case in terms of the epistemic significance of error possibilities, Martin Kusch (2009, 13–15) has also criticised, I think rightly, how high Boghossian has raised the bar here. Kusch makes this point by imagining a disagreement between himself and the W. P. Alston (1993) of *Perceiving God*. Kusch's point is that disagreement with Alston, who embraces what Kusch regards as a different epistemic system as his own, in virtue of Alston's embracing a principle Kusch calls 'Mystical Perception' (p. 9) can very well bring him to a kind of 'epistemic ambivalence' (p. 14). And in characterising this epistemic ambivalence, Kusch (2009, 14) draws from an example of the engendering of 'moral ambivalence' from Wong (2006) – an example which bears unsurprisingly close connections with the the Knobe–Nichols example. As Wong (2006, 5) puts it (in the case of moral ambivalence): 'There is...a kind of moral disagreement that poses special difficulties for universalism. This kind of disagreement evokes a complex reaction I call "moral ambivalence". We see that reasonable and knowledgeable people could have made different judgements than we are inclined to make about these conflicts, and any prior convictions we might have had about the superiority of our own judgements get shaken. Moral ambivalence is the phenomenon of coming to understand and appreciate the other side's viewpoint to the extent that our sense of the unique rightness of our own judgements get destabilized. In other words, the most disconcerting kind of moral disagreement is...a disagreement in which coming to the other side brings along an appreciation of its reasons' (cited also in Kusch (2009, 14)).
9. Note that this is the case even if the doubts engendered are not themselves rational to have. (This might, for instance, be the case if in the original position we find our interlocutor – who endorses a different set of norms – in a poor state, one who seems to not be getting on well.) As Jennifer Lackey (2008) has put it, beliefs or doubts (irrational or not) can constitute psychological defeaters in virtue of being *had* (not in virtue of being rational to have, as in the case of normative defeaters), and the recognition of a psychological defeater plausibly provides some epistemic motivation for *defeating* this defeater – in this case, perhaps by finding additional considerations that support one's own epistemic system over one's interlocutor's in the original position.
10. See Blome-Tillmann (2009) and Williams (2001) for some discussion for why such a requirement might be too strong; however, both are challenging the idea by challenging Lewis' (1996) rule of attention. Note though that the critique that this is too strong is based upon a conception of how an alternative is ruled out that is not as inclusive as Pritchard's (e.g. favouring support).

11. Though right now we are just interested in the issue of how the circularity is supposed to *arise*, it's obvious that Williams is also (as noted after the passage) telling the wider story of how we get to the familiar interim conclusion, according to which all epistemic systems are on an equal status. His further steps in the passage, spelled out, are: (iii) Any other individual (S1, S2...) who is trying to justify her own epistemic framework or system is in the same position as S; and (iv) If (i-iii), then the claims of any other epistemic frameworks, Y, Z...N, Y, Z...N have equal epistemic status as X.
12. See for example Wittgenstein (OC, §250 and §§341–3).
13. Wittgenstein himself is pessimistic about the prospects of any sort of rational progress in such a situation, as he suggests with his famous missionary–narrative conversion analogy.
14. OC §358–359.
15. We can imagine here a Westerner attempting to argue, *contra* a leader of the Azande, that the Western form of life (and its constituent hinges) has a range of desirable benefits compared to the Azande form of life. However, the Azande leader can simply point to desirable benefits of an Azande form of life, benefits that are desirable *by the lights of the Azande* form of life.
16. That is, in order to attempt to draw attention to characteristics of one's hinges that are 'better' epistemically than one's opponents, one must inevitably apply, as the background against which the evaluation is made, evaluative criteria that are constitutive of one's own form of life.
17. It is worth noting that Coliva (2010) and Pritchard (2010) both think the move from the Wittgenstenian view just sketched, to epistemic relativism, can be blocked – and interestingly, they try to snuff out the argument early on, by resisting a presupposition of the scenario just described – one where we imagined that there is 'no rational basis by which either agent could properly persuade the other to revise their view'. Thus, Coliva's and Pritchard's attempt to block a move from Wittgenstenian 'anti-foundationalism' to relativism calls upon reasons to further cast doubt upon the thought that we will have real-life cases where no Archimedean meta-norm *could* be appealed to. The Coliva/Pritchard move against the (would-be) Wittgenstenian relativist is instructive because it nicely sets the scene for why Ian Hacking's appeal to 'styles of reasoning' might have special promise. For her part, Coliva (2010), thinks that Wittgenstein's hinge epistemology motivates epistemic relativism only on what she regards as an implausible 'naturalist' reading of Wittgenstein (e.g. Strawson (1985)), where it is regarded as 'just part of our lives to take certain propositions, theories and methods of justification for granted, and thereby act in accordance with them', and that accordingly 'it is a mere accident that we act in certain ways and that our lives are what they are' (p. 2). In a similar spirit, Pritchard draws from Wittgenstein's 'river-bed' analogy to suggest it's false that different individuals can be committed to *radically* divergent hinge propositions, even if different individuals can be committed to some and even many different hinges.
18. See Carter & Gordon (2014) for a fuller discussion and critique of Hacking's strand of epistemic relativism.
19. Drawing from Alistair Crombie (1981), Hacking argues that there are many styles of reasoning that have come in and out of existence, and further, that 'We cannot reason as to whether alternative systems of reasoning are better

- or worse than ours, because the propositions to which we reason get their sense only from the method of reasoning employed. The propositions have no existence independent of the ways of reasoning towards them' (Hacking 1982, 334).
20. He adds, revealingly, that: 'What we have to learn is not what they took for true, but what they took for true or false' (For example, that mercury salve might be good for syphilis because mercury is signed by the planet Mercury which signs the marketplace, where syphilis is contracted.) (*Ibid.* 330.)
  21. See Carter & Gordon (2014) for further discussion.
  22. Examples of such frameworks Kuhn offers are the Aristotelian analysis of motion, the Ptolemaic account of planetary positions in *Almagest* and Lavoisier's application of the balance as put forward in *Traité élémentaire de chimie*. See Bird (2013 §3).
  23. During periods of what Kuhn calls normal science, when we engage in what he calls 'puzzle solving' from within a particular paradigm, we are *not* engaged in the task of questioning the very principles or distinctive concepts of the paradigm itself – just as, by analogy, on the Wittgenstenian view, when we are investigating against the background of certainties, we are not calling into question the background. A famous component of Kuhn's view is that paradigms can themselves 'shift' as a result of periods of crisis associated with the recognition of anomalies within the dominant paradigm. For Kuhn, this occurs during phases of 'revolutionary' science – when anomalies move us to replace old paradigms with new ones. To be clear, it's unlikely that Kuhn would regard an original-position-style case as an 'anomaly'. Anomalies occur from *within* a particular paradigm.
  24. See Bland (2014, 2) for discussion on this point.
  25. Thanks to Aidan McGlynn for suggesting this expression.
  26. I'm taking this characterisation of basic knowledge from Kallestrup (2012, 396). See also Cohen (2002).
  27. A proponent of BKEP might point out that even if epistemic circularity were to arise were one to have to *appeal* to – and in doing so, *apply* – one's own epistemic system in order to be justified in believing (O), a claim within one's own system – there is no reason to think that we actually have to do so, provided we are externalists. On this line, even if, when coming into contact with one who claims to find (O) 'a ridiculous principle', and we regard it as relevant to *justify* (O), our *actual* justification for (O) will not depend on our being able to do so successfully. Again, via BKEP, all that needs to be the case is that our belief that O is, in fact, the product of a reliable process.
  28. One key response, defended by W. Alston (1993), is that since scepticism is implausible, and since we simply *can't* non-circularly come to know that our faculties are reliable, epistemically circular justification that our belief forming processes are reliable can't be so bad. See Bergmann (2008, 31) for discussion and also for another line of defence, also embraced by Pryor (2000) and Schmitt (2004).
  29. The bootstrapping line of reasoning, following Vogel (2000) and Kallestrup (2012), goes as follows: Suppose one employs reliable visual perception that the table is red. By BK, one can know this even if one doesn't know that the source (perception) is reliable. Now, also a consequence of BK is that (provided S's introspection is reliable) S can also know that it was S's visual perception

- (as opposed to, say, S's hearing this by testimony) which produced S's true belief that the table is red. But, awkwardly, S can then infer deductively that S visual perception has on this occasion generated a true belief. Iterate this process a bunch of times, and S then can know that S's perception has many times produced true beliefs. Now, as Kallestrup (2012, 397) sums up the situation, having amassed track-record evidence that [S's] visual perception is veridical on these occasions, S inductively infers that such perception of hers is reliable. As induction is reliable, S arrives at knowledge of that conclusion, in short: (i) The table is red (ii) S's visual perception produced the belief that the table is red (iii) S's visual perception produced a true belief that the table is red (iv) Repeat (v) *S's visual perception is reliable*. But it seems illicit to end up at (v). Remember, BK seemed to block a straightforward kind of epistemic circularity by allowing one to know (i) *without* knowing (v). But, now it looks like BK entails that once you have (i), you can 'bootstrap' knowledge of the proposition (v) you originally didn't know.
30. For some further rationale for the more general thesis that the epistemic demands on assertion can shift across conversational contexts, see Turri (2010).
  31. See Evans-Pritchard (1937) for the classic treatment of the case of the Azande. Cf. Boghossian (2006a).
  32. One might object and assume that one is relying on an epistemic principle: *Believe what the epistemic angel says*. However, this mischaracterises the case. The consultation of the epistemic angel is not meant to be a specific case of a more general kind of principle such as: rely on the testimony of one who appears authoritative. Rather, let's envision the beliefs about the processes formed via the assistance of the epistemic angel as a *sui generis* kind of belief acquisition. As such, we should not regard this as an epistemic principle – where principles are applicable. (This *sui generis* event is not.)
  33. Fumerton (1995, 180). Cited also in Boghossian (2006a).
  34. The distinction I am making here is a broad one. For a more nuanced discussion of applying rules in a legal context, see Hage (1997, Ch. 3).
  35. Just as what explains one's course of action can 'come apart' from the rule one invokes to justify it, likewise, one can follow a rule in justifying something while claiming that one is following a different rule. For example, one might reason in accordance with modus ponens to justify modus ponens; however, one might have an implausibly rigid conception of what counts as reasoning in accordance with modus ponens (suppose: one thinks that you have to state the full modus ponens principle explicitly as a premise in the argument). In such a case, we have application circularity in that one applies modus ponens in the course of justifying modus ponens even though one claims not to be doing so.
  36. A corollary of Boghossian's point is that: if rule-circular arguments *fail* to transfer warrant from their premises to their conclusion, we should expect this result to flow as well in some natural way from the conditions that govern warrant transfer quite generally.
  37. As Moretti & Piazza (2014) note, most epistemologists who weigh in on this debate use the term 'warrant' though 'they all seem to use the term 'warrant' to refer to some kind of epistemic justification' (2014, §1) and in doing so 'broadly identify the epistemic property capable of being transmitted with

- propositional* justification'. I'll be using these terms interchangeably in the discussion in this section because (for our purposes) nothing hangs on this terminological difference.
38. Pryor's case used to illustrate Type-4 without Type-5 dependence proceeds as follows: (a) I am introspectively aware that I'm having a cold sensation now; (b) I am having a cold sensation now; (c) Therefore, I'm not making a priming mistake right now. However, it's not entirely clear from Pryor's own discussion why, exactly, such a case should work. As Ram Neta (2007, 17) remarks: 'for the example above to do the argumentative work that Pryor wants it to do, we need to know why we should believe that what makes me propositionally justified in believing the conclusion is not precisely the same thing that makes me propositionally justified in believing the premises.'
  39. For a similar move, with respect to the variety of epistemic circularity that arises in bootstrapping arguments, see Markie (2005).
  40. For example, if one antecedently doubts (O)—presumably because one believes there is some fundamental disconnect between how things appear and how they are—then it should be very odd if one somehow were to set such doubts aside and embrace inference to the best explanation. Plausibly, abductive reasoning licenses one to conclude that redness of a wall (in ordinary circumstances) best explains the visual appearance of as of a red wall. And so if one *antecedently* doubted (O) and doubted that its visually seeming to one that there is a red wall in front of one justifies one in believing *that* there is a red wall in front of one, then this at least to some extent undermines one's justification for abductive reasoning characteristic of IBE
  41. One might argue the 'pro' case here as follows: (O) is not *merely* some peripheral principle, one that Western Science could proceed without. In fact, in the absence of (O) Western Science would be unrecognisable. (O) after all is implicit in the practice of *experimentation*, where visually observed results are recorded; (O) is also implicit in how observed results are registered to inform other experiments. Whatever justifies one in accepting (O) plausibly is at least *among* the conditions that justifies me in thinking that Western Science (within which this principle is so central) is justified. And if that's the case, then the justificational structure in the line of reasoning from 24–26 is defective in that justification (or warrant) does not transmit from 24 to 26.
  42. Boghossian sets things up so as to make explicit that justifying modus ponens is necessary for justifying the corresponding epistemic transmission principle: If S is justified in believing that *p* and is justified in believing that 'If *p* then *q*', and S infers *q* from those premises, then S is *prima facie* justified in believing *q*. This principle will be true, Boghossian notes, provided MP is necessarily truth preserving. He writes: 'if S is to know that his fundamental transmission principle is true, he must, at a minimum, be justified in believing that [MP] is true. So our question about the knowability of epistemic principles becomes: Is it possible for S to be justified in believing that all arguments of the form modus ponens are necessarily truth-preserving?'
  43. We can envision, following Enoch and Schechter (2008, 576) a parallel example involving IBE, as follows: (i) We use IBE, both when doing science and in our everyday commonsensical reasoning. (ii) Our scientific and commonsensical reasoning have been tremendously successful. (iii) If IBE

- had not been reliable, this success would have been utterly mysterious. (iv) Therefore, IBE is reliable.
44. Alston (1986) offers the following view of what would be both necessary and sufficient for a belief *p* to confer warrant to another belief *q*. (A) *S* is justified in believing the premises, *p*. (B) *p* and *q* are logically related in such a way that if *p* is true, that is a good reason for supposing that *q* is at least likely to be true. (C) *S* knows, or is justified in believing that the logical relation between *p* and *q* is as specified in (B). (D) *S* infers *q* from *p* because of his belief specified in (C). Obviously, Pryor and Wright can both agree that satisfying A-D isn't (and contrary to what Alston suggests) sufficient for warrant transmitting. After all, A-D can all be satisfied in Type 5 dependence cases.
  45. See Ryle (1945, 5–6) for a helpful summary of this famous regress.
  46. For a more detailed discussion on this point, see Carter and Pritchard (forthcoming), on the issue of how to justify inference to the best explanation.
  47. Note that the argument from parity here is meant to be *inductive*, not deductive. The perceptual warrant transmission debate features two prominent anti-sceptical strategies *in the face of* epistemic circularity. We have *prima facie* reason to expect that in the face of what I argued to be the same *kind* of threat from epistemic circularity, we should expect the same options should be the salient ones. The rationale for embracing this inductive argument is defeated of course if it can be shown that the salient options in the perceptual warrant debate *should* include something like 'perceptual relativism' even though perceptual relativism is not considered a serious contender in the perceptual arena. Absent reason to think perceptual relativism should be a serious contender in the perceptual warrant debate, we've got some *prima facie* reason to expect that the same anti-sceptical options should be the salient ones in the face of this same variety of epistemic circularity when the debate turns to epistemic principles.
  48. Compare with the previous note in this chapter on blind entitlement.
  49. As Robin McKenna has pointed out to me, while the established debate about warrant transmission in the perceptual case doesn't feature relativism *per se* as a salient option, there is some precedent for embracing a contextualist semantics for perceptual warrant attributions. Even though contextualism is not a salient strategy embraced by Wright and Pryor, perhaps it should be. And if it were, then (by parity) we should expect something like a contextualist anti-sceptical strategy in the face of the same kind of circularity, at the level of epistemic principles, along with dogmatism and conservatism. (Or, perhaps: contextualist versions of dogmatism and conservatism).

## 6 Replacement Relativism: Boghossian, Kusch and Wright

1. Burnyeat's (1976) rationale seemed to be one of charity: either attribute a translation principle such as POT to the relativist or risk losing a grip on what the relativist is saying. Recall that Burnyeat's thinking here (again, in the context of assessing Protagorean subjective global relativism) went as follows: while whereas the absolute prefix 'It is (absolutely) true that' can be iterated over and over, a relativistic prefix such as "It is true for Protagoras that" ... admits of only limited reiteration. At some point, though we may not

- be able to say just where, Protagoras must stop and take a stand.' Failing to do so, he thinks, will be at the cost of 'losing grip' (p. 194) of relative truth.
2. Boghossian (2006b, 21).
  3. Note that Kusch is not regarding the position outlined in his (2010) paper as anything like 'the correct view of relativism', but rather as showing that one can carve a workable relativistic position out of the very material that Boghossian gives us – *viz.*, to show that the replacement model is not doomed as Boghossian thinks. Thanks to Martin Kusch for discussion on this point.
  4. See Boghossian (2006a, 86). Boghossian also formulates the epistemic relativist thesis this way in Ch. 5, *Epistemic Relativism Defended*.
  5. This discussion is from Boghossian (2006a, 86).
  6. In making this point, Boghossian considers an analogy. Suppose one regards all propositions of the form [Jack is immortal] as false. If so, there is 'no choice but to think that' the following [All men are immortal] is also false. But the relationship between a particular epistemic judgment and an epistemic principle is according to Boghossian's relativist like the relation between the proposition that [Jack is immortal] and [All men are immortal] (See Boghossian (2006a, 86)).
  7. Boghossian (2006a, 87). My italics.
  8. See Kusch (2010, 166).
  9. Boghossian (2006a, 88). My italics.
  10. Cf. MacFarlane (2014, 84–85) for a discussion and response to a different kind of incompleteness objection that's been raised against Kaplanian (1989) time-neutral content – *viz.*, that it is not suited to be the objects of assertions and beliefs.
  11. See Boghossian (2006b).
  12. The version of incompleteness-theoretical fictionalism which Boghossian attacks has a 'damned if you do, damned if you don't' relationship with entailment. It is because after all that error-theoretical fictionalism is a view on which the falsity of first order judgments, by entailment, generates false epistemic principles (which one must then explain how they can be accepted and endorsed) that originally made incompleteness theoretical fictionalism look to have a comparative advantage, *given that* it looks like incomplete first order principles don't entail anything. *However*, it is also problematic if there are *no* entailment relation between first order judgments and principles given that general principles are *supposed* to entail first order judgments, in so far as we are to make sense of how epistemic principles adjudicate between what's justified/unjustified at the first order.
  13. Again, we should be reading Kusch as putting forward a kind of conditional: that if we want to semantically model the epistemic relativist's core insight within a replacement model, that pursuing the strategy Kusch outlines is a workable way to do so. This is, again, weaker than a claim Kusch is not defending, which is that this is the best and only way to represent the view.
  14. Sites also in Kusch (2010, 168).
  15. Kusch defends this point at greater length later on in the paper, on p. 170, though he does not take his argument to rest on having made 'a conclusive case for this view of the ordinary person' (2010, p. 170).
  16. Kusch (2010, 169). notes that Boghossian (2006a, 83) allows for this possibility in a different context. I take it Kusch is referring to Boghossian's

pointing out what he takes to be an inadequacy with what Boghossian calls the 'traditional' refutation of epistemic relativism – one modelled off of the line taken by Nagel (2001) against the global relativist. The traditional refutation – tailored to the epistemic relativist – forces the epistemic relativist to say either that the statement of epistemic relativism is objectively justified (in which case the thesis refutes itself) or justified only relative to the relativist's epistemic system (in which case, according to the Nagel line, it's just a report of what the relativist finds agreeable to say). In critiquing Nagel's posing the 'subjectivist' horn of this dilemma on the epistemic relativist, Boghossian insists that the subjectivist horn does not commit the relativist to regarding the statement of relativism as nothing more than what the relativist finds agreeable to say. Boghossian writes: 'Indeed, it doesn't even follow that he is saying that relativism is justified only relative to epistemic principles that are *unique* to relativists. For all we are entitled to assume, he may mean that relativism is justified by a set of principles that are endorsed by relativists and non-relativists alike' (Boghossian 2006a, 83). While Kusch is technically right that Boghossian is allowing for the possibility of what Kusch is endorsing and which I'm calling the 'epistemic community thesis', it seems Boghossian is really only allowing for this as an *epistemic possibility* for individuals who are both epistemic relativists and who take the subjectivist horn. That is, it remains a live option for one who takes the subjectivist horn in Nagel's dilemma to suppose that relativism is justified by a set of principles that are endorsed by relativists and non-relativists alike. It's unclear whether Boghossian would go further to regard the epistemic community thesis as more than merely an epistemic possibility for a relativist.

17. My italics.
18. See here also Matthew Chrisman (2007, 241) for a move similar in spirit, but under the expressivist banner.
19. This of course is not to say that, if we grant Kusch the ordinary person and epistemic community theses, then there seem to be no insuperable barriers to defending a version of replacement relativism, *simpliciter*. It's just that, with these theses as collateral, Kusch should be able to combat the particular problems that Boghossian regarded as insuperable to epistemic replacement relativism defended along incompleteness lines.
20. Consider that what properties (including doxastic properties) we will attribute to the *ordinary* person can differ dramatically across possible worlds, worlds where what is in the actual world bizarre is ordinary.
21. The relevant parallel principle, in the case of epistemic relativism, will – drawing from the same core idea motivating POT – state that 'a proposition of the form 'x is F' is true (relatively) with respect to some epistemic framework/system, E, if and only if, 'x is F *relative to epistemic framework/system E*' is true (absolutely).
22. McFarlane (2014, 33) raises a similar strand of objection. In remarking on Boghossian's inclusion of the relationist clause in his formulation of relativism (2006, 56), MacFarlane writes that 'Boghossian's relativist takes a speaker who utters "snow is white" to have asserted that according to her world-theory, snow is white. But the relativist need not, and should not, hold that to put *p* forward as true for oneself is to put forward the claim *that p is true for oneself*. The point of "for oneself" is not to characterize the *content*

that is asserted, but to characterize what the relativist is *doing* in making her assertion: putting its content forward as *true for herself*.'

23. Wright (2008, §§4–5) does this though in the context of criticizing primarily the matter of whether Boghossian's own criticisms of epistemic relativism (criticisms that were targeted against epistemic relativism formulated in the replacement template) would also apply *mutatis mutandis* were the relationist clause dropped and the view tweaked so as to fit within the assessment-sensitive framework.
24. This was, in particular, a problem that raised its head *vis-à-vis* Pyrrhonian arguments for epistemic relativism and dialogic arguments for epistemic relativism. Though, the problem raised its head as well with respect to the incommensurability–circularity–relativism sequence outlined in Chapter 5.

## 7 A Different Kind of Epistemic Relativism

1. For an overview of this issue, see MacFarlane (2014, Ch. 2 and especially p. 30).
2. See, especially, MacFarlane (2014).
3. See Baghramian and Carter (2015, §5) for a more detailed discussion.
4. Cf. *non*-indexical contextualism (e.g. MacFarlane 2009 and Brogaard 2008).
5. See Rysiew (2011) for a comprehensive discussion. Cf. McKenna (2015b) for a detailed contemporary review of literature for and against this kind of view.
6. See Baghramian (2004) for a helpful survey of various kinds of relativism, many of which will be happy to include what epistemologists call contextualism as a form of relativism.
7. The idea that epistemic standings such as knowledge are always only relative to some perspective or standard features in views that have either been explicitly or tacitly acknowledged as versions of epistemic relativism, such as (among others) those found in Bloor (1976), Kuhn (1962), Hacking (1982), Feyerabend (1999), Latour and Woolgar (1986) and Rorty (1979).
8. See MacFarlane (2014, 22) for a discussion of this point, as concerns contextualist as opposed to relativist treatments of 'tasty.' MacFarlane (p. 33, fn. 5), as we might have expected, refers to Boghossian's version of relativism as a form of contextualism.
9. My italics.
10. Though note that the terminology 'context of assessment' is perhaps optional. See for example Richard (2004) who thinks that a view can come out on the interesting side of the line between relativism and absolutism provided that we allow for parameters in Kaplanian circumstances of evaluation which are not initialized by the context of use. As Robin McKenna has suggested in conversation, whether the notion of a 'context of assessment' is itself *needed* might be a matter of presentational utility rather than theoretical utility.
11. MacFarlane, (2014, 65) articulates what one must be committed to in order to *be* a relativist twice over (at different levels of generality). '**Relativism about truth**. To be a relativist about truth is to hold that languages with assessment-sensitive expressions are at least conceptually possible. This is a position one might endorse or reject on nonempirical, philosophical grounds; what it requires is that one come to understand what it would be for

an expression's extension to depend on features of the context of assessment. By contrast, relativism about truth in English is at least partly an empirical thesis: **Relativism about truth in English**. To be a relativist about truth in English (or some other natural language) is to hold that some expressions of English are assessment-sensitive'.

12. This presentation follows closely the formulation of this view articulated in Baghramian and Carter (2015, §5).
13. As MacFarlane (2014, 67) remarks "The threshold of relative truth is only crossed when we give a semantically significant role to the context of assessment.
14. Kaplan's view, specifically, was that the need for particular parameters in the circumstance of evaluation was a function of the *non-specificity* of certain propositional contents with respect to world, time and location. On Kaplan's view: 'A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well. The amount of information we require from a circumstance is linked to the degree of specificity of contents and thus to the kinds of operators in the language....' (1989: 502)
15. This is, of course, not to suggest that the proponent of the relationist clause could not (beyond what Boghossian has already said) offer some further rationale. One such rationale might be to point to the tradition of thinking of relativism as motivated by the thought that individuals are 'talking past one another.' This idea might well be better captured by Boghossian-style relativism, with the relationist clause, than by MacFarlane-style relativism. Thanks to Robin McKenna for raising this point.
16. As MacFarlane puts it, 'Boghossian models his version of truth relativism on Gilbert Harman's version of moral relativism (Harman 1975), which is essentially a form of contextualism about terms of moral evaluation' (2014, 33 fn. 5).
17. See Baghramian & Carter (2015) for an overview of some of these proposals. See also Cappelen & Hawthorne (2011) for a sustained attack against any such kind of proposal.
18. A similar view has also been defended by Richard (2014), though I'll be focusing on MacFarlane's presentation, which is more developed.
19. In some discussions (e.g. 2014, §8.4.2), expressivism is also included, but the three options listed here remain constant.
20. Note that this is not the case in MacFarlane's (2005) earliest presentation of the view, in which he opted for a generic notion of an 'epistemic standard' with no further commitments about what shape satisfying that standard might take.
21. Note that the particular details of this conundrum aren't necessary for MacFarlane to set things up as he wants to; the same kind of argument could as well be extracted from a more familiar DeRose-style (1992) bank case.
22. See, for instance, Brown (2006), Nagel (2010), Reed (2010) and Williamson (2000).
23. For example, the standards do not vary across contexts of use or across circumstances of evaluation, in the sense of Kaplan (1989).
24. This is not entirely fair for two reasons. Firstly, while it would be dogmatic to insist one does know one can rule out the error-scenario but offering no reason to think one can do so, taking a 'neo-Moorean' route out of the conundrum needn't be dogmatic in this sense. The neo-Moorean might well make

- her move on the basis of a detailed (and non-question-begging) rationale. But even more, there is a false choice at play here, one which glosses over the difference between *discriminatory* and *favouring* epistemic support (Pritchard 2010; Carter & Pritchard, *forthcoming*) in a way tacitly presupposes that the dogmatist *can't* actually rule out the alternative in question, simply because discriminatory support is not available. I'll explore in more detail in the critical discussion how this ends up being problematic for MacFarlane.
25. Cf. Non-indexical contextualism, which is a different view, and which will be discussed later.
  26. That contextualism offers such a resource for evading the sceptical problem has been taken by contextualists as a major point of recommendation for the view. See for example DeRose (1995) and Lewis (1996).
  27. MacFarlane suggests on p. 177, fn. 2, that he is not taking a stand on what it is to rule out an alternative. However, he betrays in his discussion on pp. 177–178 that he has a rather robust idea of ruling out an alternative in mind – one where the ruling out involves *actions* (e.g. checking one's pockets, etc.). On at least one standard way of thinking about ruling out alternatives (e.g. Pritchard (2010)), possession of abilities or possession of certain kinds of favouring support can entitle one to rationally dismiss an alternative. I'll revisit this point later in this chapter.
  28. Cf. Brown (2006), Buckwalter (2010), Hawthorne (2004) and Stanley (2005) for some arguments against. Cf. McKenna (2015) for a defence of contextualism against arguments to the effect that it cannot explain variability data.
  29. See DeRose (2009). However, see McKenna (2014) for a recent attempt to bypass this problem.
  30. See Williamson (2005).
  31. See MacFarlane (2014, 177) for a formal representation of the differences.
  32. *Ibid.*, p. 182.
  33. Cf. Marques (2015) for a recent challenge to MacFarlane's appeal to retraction. See also Stanley (*forthcoming*) for a critique of MacFarlane's dismissal of subject-sensitive invariantism.
  34. *Ibid.*, pp. 184–185. While MacFarlane engages briefly (e.g. pp. 185–187) with some attempts by Hawthorne (2004) and Stanley (2005) to soften the force of these worries, he is not convinced.
  35. It is also worth pointing out that contextualists, SSI-ists and insensitive invariantists argue amongst themselves about who best accommodates the variability data, and so the comparative 'costs and benefits' approach is not unique to the truth-relativist.
  36. Cf. Stanley (*forthcoming*).
  37. MacFarlane terms *postsemantics* the definition of truth at a context in terms of the *semantics proper* – viz., the definition of truth at a context and index (MacFarlane 2014, 59).
  38. Carter & Pritchard (2015) builds upon, and develops further, the 'two-tiered' relevant alternatives approach originally put forward in Pritchard (2010).
  39. For related discussion on this point, see Pritchard & Kallestrup (2012).
  40. The *locus classicus* is Gettier (1963). See also Pritchard (2005) for detailed discussion.
  41. Put differently, the target belief is, in Pritchard's (2005) sense *veritically lucky*: that is, given the initial conditions of the belief's formation, there are near-by

worlds where the belief is incorrect. The distinction between intervening and environmental luck is Pritchard's (2005). See Pritchard (2007) for further elucidation.

42. This is a spin-off of Chisholm's (1977) 'sheep in a field' case, a case of what Pritchard (2005) calls 'intervening' epistemic luck.
43. I revisit this point again in more detail later in the chapter.
44. MacFarlane (2014, 189).
45. As described, the kind of case suggested here looks as though it might be directed toward contextualism just as well as against MacFarlane, and if so, it might seem as though Case 3\*, targeted at the relativist, attempts to prove too much. However, the contextualist has a way to at least in principle deal with cases like this that is not equally available to the relativist. Lewis (1996), in his famous paper 'Elusive Knowledge', has a number of 'rules' (e.g. the rule of belief, the rule of resemblance) which are designed to assist the contextualist in handling cases like 3\*. Consider, for example, Lewis's 'rule of actuality' and 'rule of resemblance'. The rule of actuality says the possibility that actually obtains for the subject are always relevant; and the rule of resemblance says that any possibility that 'saliently' resembles a relevant possibility is relevant. Lewis thinks these rules, working together can deal with barn façade cases, as he mentions on p. 557. He writes, 'Unbeknownst to me, I am travelling in the land of the bogus barns; but my eye falls on one of the few real ones. I don't know that I am seeing a barn, because I may not properly ignore the possibility that I am seeing yet another of the abundant bogus barns. This possibility *saliently resembles actuality* in respect of the abundance of bogus barns, and the scarcity of real ones, hereabouts' (Lewis 1996, 557, my italics). There is nothing *in principle* illicit about 'importing' objective features of the 'subject's situation' (to use a term from Robin McKenna) into the context of use, and so for the contextualist, the primary challenge (for handling 3\*-style cases) will simply be to capture, in the right kind of way, *which* objective features of the subject's situation the context of use should be thought of as including. But for the relativist, things are much more complicated. Just suppose that MacFarlane were to try to handle cases like 3\* by helping himself, *a la* Lewis, to similar sorts of rules. On this supposition, the ensuing problem is really twofold, the first problem concrete, the second more abstract. Firstly, the concrete problem is that MacFarlane himself seems to cut such an option off, in order to avoid the temporal and modal embedding objections that faced SSI. In remarking (e.g. 2014, 187–8) about desiderata his relativist view must satisfy to avoid what he regards pitfalls of SSI he stipulates that his view has the benefits of preserving that "the way in which the alternatives a subject must rule out in order to count as "knowing" vary with context. However, it would not...*join SSI in taking this variation to be keyed to the circumstances of the subject to whom knowledge is ascribed*, since that is what makes it difficult for SSI to explain our judgments about... [sic. temporal and modal] embedded cases...' (MacFarlane 2014, 188). If keying relevant alternatives to the context of assessment positively *excludes* that the relevance of alternatives is keyed to the circumstance of the subject to whom knowledge is ascribed, then MacFarlane couldn't by his own lights advert to Lewis-style rules (such as actuality and resemblance) in order to handle cases like 3\* insofar as such appeals violate the constraint he notes at (2014,

- 188) meant to avoid what he regards as a pitfall of SSI. The more abstract point, though, which makes this kind of move a difficult one to suppose MacFarlane could help himself to as the Lewis-style contextualist does, is this: it is the very fact that features of the context of assessment, *rather than* various other features of contexts (e.g. context of use, circumstance of subject of attribution) is what's determining relevance of alternatives that distinguishes his view as properly *relativistic* by his own lights. Contextualists after all regard that truth values of knowledge ascriptions are *absolute*. To the extent that the relativist imports objective features of the subject's situation into the *context of assessment* (in order to deal with 3\*-style cases) we quickly lose our grip on how the view's supposed to remain properly relativistic, such that (as Wright (2007, 262) puts it in describing the key thrust of 'new age' relativism), 'the truth-value of the utterance can vary, even though the context of its making and the associated state of the world remain fixed.' Thanks to Robin McKenna for prompting discussion on this point.
46. My reason for this qualification will be made clear soon.
  47. However, as I'll discuss later, there is a second strand of secondary relevant alternatives with respect to which MacFarlane's view does much worse. The kind of secondary relevant alternatives we've engaged with so far are those which become secondary relevant in virtue of an agent's becoming aware of them. A second strand, discussed in detail in Carter & Pritchard (2015), is secondary relevant not because one becomes aware of the alternative in question but because one *should* have become aware of it.
  48. Though, to reiterate, I don't think we should, once the *normative* variety of secondary relevance is considered. See fn. 29.
  49. Thanks to Brian Rabern for helpful discussion.
  50. That holograms chaffinches continue to swarm around Alan, along with the real ones, no matter what things are like in the environment from which Charles is evaluating Liz's claim about Alan, means that there are inevitably near-by worlds where Alan is believing falsely that he's looking at a chaffinch, regardless of what things are like across context of assessment in which Liz's claim about Alan is being assessed.
  51. Originally proposed by Ginet (1975). See Goldman (1979) for the canonical expression of the point. While I regard, following here Pritchard (2009a, 2009b, chs. 3–4; 2009c ; 2012a) and Haddock, Millar & Pritchard (2010, chs. 2–4) the anti-luck insight as a *fundamental* platitude that should inform our theorising about knowledge, it's also important to note that this viewpoint is not free from criticism. We can divide this criticism into roughly two categories of relevance, *vis-à-vis* my critique of MacFarlane. Firstly, there are those who reject the anti-luck insight, *tout court* (e.g. Hetherington (2012) and Baumann (2014), *cf.* Madison (2011)). Interestingly, though, not everyone who rejects the claim that environmental luck is incompatible with knowledge does so by challenging the anti-luck insight, construed broadly, as Hetherington and Baumann have. For example, Colaço, Buckwalter, Stich, & Machery (2014) have challenged this incompatibility of *specifically* environmental luck with knowledge on the basis of experimental philosophy results (see also Lycan (2006), Turri (2011) and Turri, Buckwalter & Blouw (2014)) and similar experimental results have been replicated more recently by Turri, Buckwalter & Blouw (2014, 386). Some representative statements

- of the view that knowledge excludes barn-façade style environmental luck include: Swain (1978), Lewis (1996), Kvanvig (2003, 2004), Pritchard (2003, 2004, 2005, 2007, 2009), Steup (2008), Luper (2010), Madison (2011), Kelp (2012), Jarvis (2012), Carter (2010, 2011, 2013) and Carter, Jarvis & Rubin (2013a, 2013b, 2014), Carter & Pritchard (2013, 2014a, 2014b).
52. For expressions of this idea, see Pritchard (2005; 2007) and Carter (2009; *forthcoming*).
  53. Thanks to Robin McKenna for suggesting this way of presenting the point.
  54. This is not to become aware of some proposition  $p$  such that if  $p$  is true the target proposition is false. See Schaffer (2001).
  55. Carter & Pritchard (2015).
  56. This case is adapted from Carter & Pritchard (2015, 15).
  57. One rationale for this, which I won't take up in detail here but which I find plausible, is found in Baehr's (2009) criticism of evidentialism. On Baehr's line epistemic justifiedness with respect to a proposition at a time has to be more than just a matter of the evidence one in fact possesses at  $t$ . One might after all fail to possess other evidence through epistemically vicious inquiry. The wider point is that facts about epistemically irresponsibility (e.g. lazily overlooking evidence) can bear on how we should epistemically appraise a belief.
  58. For some helpful discussions of epistemic defeat, see especially Bergmann (1997), Lackey (1999, 2005, 2010) and Pollock (1986) See also Sudduth (2008) for a contemporary overview.
  59. This is essentially the view embraced by W. Clifford (1879/1999), in his famous dispute with William James.
  60. For one thing, as Goldberg notes, there is intuitively *something* that's gone wrong in the kinds of cases he's highlighted. The proponent of Extreme Cliffordism who wants to preserve this insight is going to have to say that these are cases of a *moral* or *practical* wrong. Combine this with the Extreme Cliffordist's commitment to thinking that the agents in these cases are *epistemically justified*, and what results, Goldberg says, is that 'the proponent of [sic. Extreme Cliffordism ] is committed to thinking that the moral or practical wrong committed by the subjects above have no bearing on the epistemic status of her belief. But do we really want to say e.g. that, while the physicist's belief (to the effect his theory is empirically well-supported) is justified, the physicist is morally wrong to ignore the comments to the contrary by his colleagues?" (*ibid.*, §5). Even if I were inclined to Extreme Cliffordism already (which I'm not, for broadly reasons of the sort highlighted by Baehr (2009)), I'd regard this consequence which Goldberg's highlighted as a seriously troubling one.
  61. This is as the view is outlined in Carter & Pritchard (2015).
  62. That is, we can simply think of *what matters vis-à-vis* relevance as fixed by what the individual in the context of assessment is aware of. Obviously, there remains the independent issue of whether keying awareness secondary relevant alternatives to the data really does best explain our patterns of knowledge attribution better than contributor positions. The point of Case 2\* was not to adjudicate this point, but just to establish that (unlike with the case of primary relevant alternatives) the way that these alternatives become relevant could come about as a matter of how things stand at the context of assessment.

63. This is intentionally generically formulated; proponents of RVE fill in the 'because' relation differently. Greco (2010) has opted for 'causal-explanatory salience' in earlier work, (language that has been adopted also by Pritchard (e.g. 2010, 2012) who defends a weaker version of virtue epistemology than RVE). Greco (2012) has modified this view somewhat more recently. Sosa (2007), Turri (2011) and Kelp (2012) by contrast opt for the language of manifestation in order to characterize the 'because' relation. See Carter (2014) for a modal proposal.
64. This is a modification of a case by Lackey (2007) aimed at robust versions of virtue epistemology. I've argued both sides of the point (e.g. Carter 2011 and Carter 2014 represent my previous view, and Carter & Pritchard 2014, 2015 represent my more recent view), on the matter of whether testimony cases are genuine *counterexamples* to RVE. More recently, I'm of the view that they are. But that's not the issue here. Remember, we're using RVE here as a toy example to show that the fact that MacFarlane's view makes bad predictions can be established even if the view drops relevant alternatives talk from MacFarlane (2014) and retreats to generic 'epistemic standard' language (a la 2005).
65. This view is established most clearly in the case of 'epistemic twin earth' cases. See for example Pritchard & Kallestrup (2011, 2012, 2013), Carter & Kallestrup (2014), Carter & Pritchard (2014).
66. See my endnote earlier, which engages with an issue raised by Robin McKenna, for clarification on why the incompatibility with epistemic anti-individualism needn't be an undoing of contextualism.

## 8 New Relativism: Epistemic Aftermath

1. And, with reference to the puzzle MacFarlane outlined, the insensitive invariantist will be stuck either rejecting closure, embracing 'dogmatism' or scepticism. Though this part is more controversial. What's uncontroversial is that insensitive invariantism can explain the variability of our willingness to attribute knowledge only with at least some special pleading.
2. Richard (2004) and Williamson (2005) along with MacFarlane (in many places) advance this argument. Relatedly, though, MacFarlane's view is that contextualism can't make sense either our evaluation of others' attributions of knowledge, or our earlier ones, are true. (MacFarlane 2014, 187). Another influential criticism of contextualism is Stanley's (2005, Ch. 2) which claims that the analogy contextualists like to make to 'tall' (e.g. DeRose 2005 and Cohen 1999) reveals a problem for the view, which is that tall is a 'gradable adjective', but 'knows' is not. See McKenna (2015) for a helpful discussion and also a reply.
3. Though see Hawthorne's (2004) and Stanley's (2005) attempts to reply to this problem by suggesting it is no worse than problems that face contextualism. Though see MacFarlane's (2014) for a criticism of their approaches.
4. Though, to reiterate, talk of epistemic luck wasn't *needed* in order to make the point.
5. Anti-individualism gains support from a phenomenon Pritchard & Kallestrup call *epistemic dependence*, as when an agent manifests very little cognitive

agency (i.e., much less than would normally suffice for knowledge), but where her cognitive success amounts to knowledge nonetheless because of factors external to her cognitive agency – as when an agent gains knowledge easily by simply trusting an informant in a friendly environment. Pritchard & Kallestrup call this *positive epistemic dependence*. *Negative epistemic dependence*, in contrast, is when an agent manifests a high level of cognitive agency (i.e., of a level that would ordinarily easily suffice for knowledge), but where the cognitive success does not amount to knowledge because of factors external to her cognitive agency. Compare for instance a barn façade case with counterpart case where the agent's local environment is friendly.

6. Perhaps, *a la* Chang (2002), there is no answer to this question; the problems might be on a par.
7. The 'ordinary person' on Kusch's view, is non-philosophical, in the sense that, the ordinary person uses 'knows' but does not have reflective views about the nature of knowledge.
8. This remark anticipates a point I'll engage with later in this chapter, one that draws from insights from Allan Hazlett (2010).
9. And, in the case of the incommensurability–circularity–relativism sequence explored in Chapter 5, other non-relativistic non-sceptical alternatives which I called (modeling off the language from the perceptual warrant debate) dogmatism and conservatism.
10. The Pyrrhonian argument for epistemic relativism, anticipated by Sankey and evaluated in Chapter 3, reached a very similar kind of gambit point. The 'end game' for a proponent of the incommensurability–circularity–relativism sequence looked a bit different. The way I set things up in Chapter 5, those who attempt to reach epistemic relativism via this sequence of arguments were left with no reason to prefer relativism to what I called 'dogmatist', 'conservative' or 'sceptical' alternatives – a choice that was framed after drawing some parallels between the kinds of epistemic circularity (what I called application circularity) that materialized off the back of incommensurability.
11. MacFarlane (2014, 176).
12. (MacFarlane 2009, 16). See also Carter (2014) for a previous attempt to approach this question, though, previously with different objectives in mind.
13. See McGlynn (2014) for a thorough critique.
14. See Ichikawa (2013) for a discussion of a contextualist treatment of  $E=K$ , one on which, 'in any given context, "S knows H" and "S's evidence includes H" will have the same truth conditions' (*ibid.*, 286).
15. While there are some senses of 'reductive' in which Stanley's view is not a *reduction*, the language 'reductive' intellectualism is useful for distinguishing his kind of view from Bengson & Moffett's version of intellectualism, where the latter denies propositionalism and so denies that knowing-how is a kind of knowing-that even though knowing how is grounded in propositional attitudes. See Carter & Pritchard (2015) for discussion on this terminological point.
16. *Cf.* Bengson & Moffett (2011) for a defence of *non-reductive* intellectualism. On their proposal, knowing how is grounded in, but not reducible to, propositional attitudes. Bengson & Moffett thus show that intellectualism and *propositionalism* can come apart. One can be an intellectualist in virtue of

- holding the view that when one knows how to do something, this is in virtue of (i.e. grounded in) one's propositional attitudes. Intellectualists can optionally embrace propositionalism, as Stanley does: a view about the *nature* rather than the grounds of, knowledge-how: namely, that the nature of knowledge-how is a propositional attitude relation.
17. In this passage, Stanley (2010, 209) is summarizing the key line defended by Stanley & Williamson (2001). Stanley's newer (2010, 2011) view embraces the same reductive model but receives a more sophisticated and sustained defence.
  18. My italics. Cited also in Grimm (2014).
  19. For explicit defences of this view, see for example, Salmon (1984), Woodward (2003), Strevens (2008), and Greco (2010). Grimm (2014) also defends a version of this view, though on Grimm's view, the kind of knowledge of causes is non-propositional, a kind of knowledge-how. What Grimm is calling the traditional view of understanding regards the knowledge of causes at issue to be propositional. For a notable criticism of the claim that propositional knowledge is sufficient for understanding why, see Hills (2010). See also Pritchard (2010), Carter & Pritchard (2014, 2015) for a line of objection to the effect that understanding why something the case involves a kind of cognitive achievement, where cognitive achievement can come apart (in both directions) with propositional knowledge.
  20. See Ichikawa & Steup (2014, §0).
  21. One line of resistance to this argument would be to suggest that the relativity of knowledge would follow from (in the sense that: it would be 'inherited' from) the relativity of *any one part* of knowledge, where that part needn't be justification, specifically. If this line were tractable, then contrary to what I'm suggesting here, the relativity of knowledge wouldn't *entail* the relativity of justification. However, I don't think this potential line of response is promising because it allows that, for some factorable property that is itself assessment-sensitive, some of its factors can be assessment-sensitive while others are not. But this principle seems open to counterexample.
  22. Boghossian (1989, 154) attributes this characterization to O'Connor & Carr (1982, 75) in the context of noting internalists' dissatisfaction with clairvoyance-style cases that would appear to qualify as cases of justification on flat-footed reliabilist proposals. For an explicit defence of this condition, see Fumerton, 1995, 183–224.
  23. *Ibid.*, 154. See in particular claim (4).
  24. For instance, perhaps one must *understand why* the evidence supports the target proposition.
  25. For an explicit statement of this form, see Greco (2013), who articulates RVE as a 'genus-species' claim.
  26. See here, for instance, Adler (2004), Baehr (2011), Battaly (2004), Carter & Gordon (2014), Riggs (2010), Roberts & Wood (2007).
  27. Cf. Baehr (2011) for a different kind of model, one on which T is an intellectual virtue in so far as it contributes to the intellectual personal worth of the agent – a view modeled off of Hurka's account.
  28. We might articulate the relevant *connection* C to E in various ways: for example, we might say that T is an intellectual virtue only if it is reliably conductive of E. Alternatively, we might say T is an intellectual virtue only

if agents, when manifesting T, are characteristically motivated toward E. Or perhaps we will follow a neo-Aristotelian model according to which virtues have a kind of 'reliable success' as well as a motivational kind of structure, and so we might insist that T is an intellectual virtue only if *both* reliability as well as motivational connections hold.

29. See Carter & Gordon (2014) for a sustained engagement with this point.
30. See, for example, Gelfert (2010) Kappell (2010) Kelp (2011; 2014), Millar (2010), Williamson (2000). One line of argument for this view claims support from the thesis that knowledge is the aim, or goal, of inquiry. A succinct way to put this idea is Millar's (2011, 63) remark: 'Suppose that you are inquiring into the matter of whether something is so. Your aim is to find out whether it is so. Since finding out is nothing less than coming to know, what you aim for is knowledge.' For some recent alternative approaches, see for example Ahlstrom-Vij & Grimm (2013), Kvanvig (2013), and Pritchard (2012).
31. For an excellent contemporary overview of the state of debate concerning knowledge norms, see Benton (2014).
32. For defences of this view, see for example Williamson (1996; 2000), Stanley (2005), Hawthorne (2004), DeRose (1995).
33. See for example Adler (2002), Hindriks (2007), Sutton (2007), Bach (2008), Engel (2004), Peacocke (1999). Cf. McGlynn (2013) for a recent criticism.
34. Riggs (2008).
35. Note that Kvanvig (2003) himself outlines this adequacy requirement as part of a wider argument to the effect that knowledge is ultimately not as valuable as we've thought, and that it is understanding (e.g. 2003, Ch. 8), not knowledge, that has the distinctive epistemic value that epistemologists have traditionally and mistakenly attributed to knowledge, and which is in part responsible for epistemology's 'myopic' focus on knowledge. That said, Kvanvig maintains that *if* one takes the orthodox line, rather than Kvanvig's own revisionary line, then it is incumbent upon one to vindicate one's analysis of knowledge at least logically compatible with the insight that knowledge is more valuable than mere true belief that falls short of knowledge. Part of Kvanvig's rationale is that any such attempt will ultimately succumb to the 'swamping problem' (see fn. 29).
36. See for example Goldman & Olsson (2009) and Olsson (2007).
37. See also Kvanvig (2009) for a more recent discussion of the problem, and also Pritchard (2011) for a novel presentation of the premises of the argument. For a recent criticism of some of the assumptions behind this problem, see Carter & Jarvis (2012). See however Dutant (2013) for a reply.
38. As things stand, robust virtue epistemologists offer what is probably the most compelling story for how knowledge – as they define it as type identical with cognitive achievement – is vindicated as more valuable than mere true belief. Roughly, the analogy is this: just as a success that is an achievement is more valuable than either a skilful failure or a lucky success that is not properly due to ability, so, getting to the truth in a way that constitutes an achievement, a hallmark of robust virtue epistemology, is valuable in a way that mere true belief that is not an achievement (e.g. not primarily due to ability) is not.
39. I've resisted this kind of thinking in previous work (e.g. Axtell & Carter 2008; Carter & Gordon 2014) and have in doing so opted for a more pluralistic

- approach to thinking about how to characterize the aims of epistemology. It remains, though, that thinking of epistemology as the theory of knowledge, its nature and scope, rules the day.
40. The presumption of this thesis probably traces to the obvious place, the rise of epistemological contextualism in the 1990s with the work of Lewis and especially Cohen and DeRose.
  41. I say indirect because, following Grice, there are sometimes *pragmatic* explanations for why we use sentences and words as we do, other than to say things that are literally true. Further, even when we do aim to say things that are literally true, we sometimes make mistakes. See MacFarlane (2005) for a nice discussion of this point.
  42. Hazlett (2010, 498).
  43. Cf. Turri (2011) for a criticism of Hazlett's defence of this premise.
  44. See, especially, Hazlett (2010, §2).
  45. Turri (2011) thinks Hazlett might be too quick here. Turri's (2011, 145) writes: 'Suppose that the philosopher's concept of knowledge is factive and the ordinary concept non-factive. For all Hazlett says, this could be the only difference between them, in which case the ordinary concept would so greatly resemble the philosopher's concept that epistemologists would do well to study it carefully. Epistemologists could profitably use intuitions about the ordinary concept to constrain theorizing about the philosophical concept. In short, a non-factive concept of knowledge could still be epistemologically interesting'. I have two remarks to make about this point. Firstly, though Hazlett did not say much in defence of (1) in the paper, there remain plenty of ways one might argue compellingly for (1). As Hazlett has pointed out to me in conversation, one might suggest that insofar as epistemologists regard knowledge as epistemically valuable, we should understand this value as at least partly derivative of the value of truth. If knowledge were non-factive, it would be perplexing why epistemologists have prized knowledge over states that fall short. Secondly, Turri's remarks above become somewhat odd; particularly, I would be inclined to think that a non-factive concept will, in virtue of being non-factive, differ so *significantly* from the philosopher/epistemologist's concept of knowledge that's hard to see how we could profitably use intuitions about the a non-factive concept to constrain theorising about the philosophical concept of knowledge, even if that non-factive concept had other commonalities with the epistemologist's concept of knowledge.
  46. As I argued in Chapter 1, the factivity of knowledge is clearly a pragmatic presupposition of both parties to perennial disputes in epistemology. Recall, again, the debate between Moore and the sceptic. The *shape* this debate takes is one that reveals factivity to be a pragmatic presupposition (in the sense of Stalnaker and Grice) taken as common ground by both Moore and his sceptical opponent. In particular, I suggested that without a commitment to thinking knowledge is factive in the shared background, we'd expect the disagreement between Moore and the sceptic to reflect what would be a *much weaker position for the sceptic*. That is, we'd have a hard time making sense of the sceptic's challenging Moore's claim to know his first premise. Moore and the sceptic *reveal*, in practice, that a truth condition on knowledge is common ground between them. Their debate would take a *very different shape* without it.

47. To the extent that this is right, we have further reason to resist the kind of line Turri (2011, 145) presses against (1) of Hazlett's argument.<sup>48</sup> Thanks to Allan Hazlett for helpful discussion on issues connected to the central argument of this chapter.

## 9 Metaepistemology and Relativism

1. In particular, what I called in Chapter 5 'dogmatist' and 'conservativist' alternatives.
2. See §8.2 entitled "Epistemic Aftermath."
3. This language is meant to be neutral with respect to the two central kinds of pictures of relative epistemological facts surveyed, the *replacement*-style picture of relative epistemic facts as absolute facts with an explicit relational clause specifying the relativity to some parameter (e.g. a la Boghossian 2006) and the assessment-sensitive model of relative epistemic facts, according to which the truth to which non-relativised statements aspire is relative truth.
4. In particular, the examples I relied on in making the point were Goldman versus Feldman about the nature of epistemic justification, and Moore versus the sceptic.
5. This was, recall, a way of thinking that aligns with Miller's (2012) 'two-component' view of generic realism, for any domain of discourse. Again, while I think Miller provides a helpful and intuitive framework for thinking about realist commitments, I appreciate that this is controversial. It would be a substantial project in its own right to positively defend one way of thinking about realism over all competitors, and I can't attempt to do that here.
6. See Chapter 1 for a discussion of the two kinds of ways meta-epistemological commitments can be incurred: by articulation (articulated meta-epistemological commitments) and by pragmatic presupposition (revealed metaepistemological commitments).
7. It would be nice if there were a general formula for distinguishing trivial as opposed to non-trivial dependencies, but there isn't. We might say, as Miller (2012) does, that the mind-dependencies that make the difference between denying and embracing the independence leg of realism cannot be *mundane* dependencies. Though, as Jenkins (2005) remarks 'it is not straightforward to say exactly which kinds of dependence are mundane.' Unsurprisingly, determining whether a meta-epistemological thesis is anti-realist in virtue of denying the independence leg of realism, by affirming that the epistemological facts that exist are non-trivially mind-dependent, requires some delicacy.
8. Downing (2013, §0).
9. Dummett's semantic proposal is one on which what facts are regarded as epistemically constrained. Likewise, Nelson Goodman's (1975; 1978) radical constructivist rejection of the claim that there is some unique mind-independent world is a paradigmatic version of 'anti-independence anti-realism'. For Goodman, there are world-facts, but the states of affairs that constitute them are *non-trivially* mind-dependent, they hold in virtue of our ways of worldmaking.
10. See Miller (2012, §§6–7) for helpful discussion.

11. As Downing remarks, commenting on a passage from Berkeley's *Principles*, 'Berkeley does not deny the existence of ordinary objects such as stones, trees, books, and apples. On the contrary, as was indicated above, he holds that only an immaterialist account of such objects can avoid skepticism about their existence and nature.'
12. Note that I think that the reasonableness of constraining what counts as an epistemic fact by appealing to platitudes does not depend on any further sympathy to Cuneo's (2007) wider realism, which he defends in both epistemology and ethics.
13. As Cuneo puts it '... content platitudes with respect to representational entities such as the propositional attitudes display one or another epistemic merit (or positive epistemic status) such as *being a case of knowledge, being warranted, being an instance of understanding, insight or wisdom* and the like, only insofar as they are representative in some respect. That is, these entities display such merits only insofar as they *represent reality aright*' Cuneo (2007, 57). See, however, the relevant qualification *vis-à-vis* content platitudes made in Chapter 1, in response to worries raised by Matthew Chrisman concerning (among other things) epistemic modals.
14. Nor does the error theorist want to deny what Olson calls the 'immanent' norms that imply hypothetical reasons. The error theorist denies transcendent norms and categorical reasons. See Chapter 1 for more discussion.
15. I should flag that the kind of hypothetical reasons that the epistemic error theorist can accept have been drawn attention to, primarily, in order to diffuse the kind of objection that an error theorist about epistemic reasons can anticipate, and which has been explored by (among others) Streumer (2012) & Shah (2008). The objection is that if epistemic error theory is true, we would have no reason to believe it. Drawing a careful distinction between the kind of reasons to believe which the error theorist is denying from those she can accept is one natural way, and one Olson (2009) pursues, to disarm this kind of threat. However, what is salvaged is probably best described as practical reasons for belief such that one might have a practical reason to believe the error theory.
16. For instance, take Kusch's (2010) example: "Otto's belief in ghosts is unjustified" – that's something that, if true, we'd ordinarily regard to be an *epistemic fact*. Now, suppose a theory T claims to preserve epistemic facts because, on that theory, some claims of *that form* are taken to be true. The import of the authority platitude is this: if, on theory T, what makes Otto's belief in ghosts count as unjustified is (for instance) facts about his *desires* – for example, that Otto doesn't like the idea of ghosts – then T has not preserved an *epistemic fact*, simply by preserving that "Otto's belief in ghosts is unjustified" can come out true on the theory.
17. For example, consider here Richard Feldman's (2000, 682) remarks on the nature of the *epistemic* ought, which features in the evidentialist's thesis that one ought to believe in accordance with one's evidence. He writes: 'If the oughts in question are supposed to be means to goals that people actually have, then it seems that only people who do have the epistemic goals just mentioned would be subject to the relevant epistemic requirements. However [the claim that one ought to believe in accordance with one's evidence] is not restricted in that way. It says that all people epistemically ought to follow

- their evidence, not just those who have adopted some specifically epistemic goals' (p. 682 my italics).
18. As I noted in Chapter 1, it's also not obvious that the content platitude would apply here, but I'm setting this point aside presently as it's the authority platitude that's of particular interest *vis-à-vis* epistemic relativism.
  19. One way to put the puzzle is like this: when Goldman Berkeley-knows *p*, this is because some mental state of affairs accessible to Goldman holds. But how could this Berkeley-accessible mental state of affairs be authoritative in Feldman-mental-space, the mental state of affairs, accessible only to Feldman within which Feldman is operating? Perhaps there are other ways of thinking about this, but any coherent story would seem to require special pleading each step of the way.
  20. Though I think there could be reasonable objections to this point.
  21. And if we have good reason to think that the authority platitude circumscribing epistemic facthood is to be unpacked on relativist lines (with intra-perspectival categoricity), then the *mind-dependence* is not interesting but trivial – as it is just part and parcel with the kind of mind-dependence that would be necessary to support facts whose authority has intra-perspectival categoricity. If we have good reason to think the relativist offers a *compelling* picture of what she's calling an epistemic fact is, then, by her own lights we are not in a good position to suggest that the relativist is an anti-realist.
  22. Is there *a priori* knowledge? Is safety necessary for knowledge? Is attributor contextualism more plausible than subject-sensitive invariantism? Is anti-luck virtue epistemology true? Is knowledge 'first'? Is epistemological disjunctivism true? Is the value of knowledge 'swamped' by the value of true belief?
  23. See H.O. Mounce (2007, 121). Note that in discussing Wittgenstein's remarks in OC §248 and §152, Mounce is careful to suggest that this quote is prefaced by Wittgenstein with the qualifier 'as one might say'; Mounce's reading of Wittgenstein is *not* as an 'anti-foundationalist' as he is often interpreted. My use of the analogy in the above does not turn on this point, however.
  24. Pritchard (2011; 2015a, 2015b) has, in several recent works, drawn reference to Wittgenstein's (OC, §§95–7) river bed analogy in order to highlight an oft-overlooked point in Wittgenstein's epistemology which is that, despite the kind of foundational role hinges can play, in virtue of which they are at any given time 'immovable', the status of hinge propositions as hinges can shift over time. The relevant passages in *On Certainty* which speak to this suggestion are: 'The propositions describing this world picture [i.e., the hinges] might be part of a kind of mythology. And their role is like that of rules of a game; and the game can be learned purely practically, without learning any explicit rules. It might be imagined that some propositions, of the form of empirical propositions, were hardened and functioned as channels for such empirical propositions as were not hardened but fluid; and that this relation altered with time, in that fluid propositions hardened and hard ones became fluid. The mythology may change back into a state of flux, the river-bed of thoughts may shift. But I distinguish between the movements of the waters on the river-bed and the shift of the bed itself; though there is not a sharp division of the one from the other' (OC, §§95–7).

25. Just as Pritchard (2012), with reference to Wittgenstein's river-bed analogy, allows for the in-principle shiftability of Wittgensteinian hinge propositions over time, he reads the Wittgensteinian as committed to the view that at least one hinge – the *über* hinge – viz., that we are not being radically deceived – is not shiftable and must be held fast. Regarding what individual hinge propositions have in common, Pritchard (2012, §4) writes that 'they each reflect what we might call an *über* hinge commitment on the part of all subjects, regardless of who they are or their personal circumstances – viz., that one is not radically and fundamentally mistaken in one's beliefs. For in all cases a mistake on the part of the subject regarding the target proposition would entail massive and fundamental error in the subject's beliefs. The *über* hinge commitment thus entails a hinge commitment to these specific propositions as well.' Along broadly analogous lines, one might suggest that *shiftable* hinges within epistemological theorising entail a commitment to something like an 'unshiftable' hinge, within epistemology – a (for lack of a better way to describe it) *quasi-über-hinge*. And if this is right, then contrary to the suggestion that what mainstream epistemological projects take for granted could 'shift' over time until something like relative epistemic facthood is taken for granted, there would be certain things (or some thing) that epistemology must take for granted and which isn't up for grabs. There's at least two non-arbitrary ways of suggesting what such a hinge might involve. Firstly, one might attempt to key such a hinge to the *content* platitude and to suggest that epistemic merit and demerit is connected to, following Cuneo, representing reality aright. On the line I'm envisaging, one might deny that *this* is something that (on broadly Wittgensteinian lines) could be rationally doubted from *within* epistemology. Alternatively, one might suggest – and here I am referencing Kusch's (2010) discussion of the distinction between *eliminating* a concept and *replacing* that concept with a successor concept, that epistemology-cum-relativist-in-the-background might be argued to be not *epistemology*, and rather, something else – in which case, as this line would go, an institution (even if over time) of relativist meta-epistemological commitments would generate debates that are no longer *a kind of* epistemology (cf. Carnap's (1930) proposed 'overcoming' of metaphysics, by simply eliminating what was done under that name). I am not attempting to defend here either of these lines, just noting that both are available to one who is resistant to the suggestion that the hinges in the background of mainstream epistemology could drift so far that they in time become relativist ones.
26. This was, to be clear, a line of argument styled off of Hazlett's (2010) 'divorce argument' which was itself made not with relativism in mind but with reference to the suggestion that the ordinary concept of knowledge is non-factive and the epistemologist's concept cannot be.
27. To be clear, the thrust of the envisaged second-wave epistemic relativist's point is that, if we attempt (as I did toward the end of Chapter 8) to 'disarm' new epistemic relativism by suggesting that the more compelling the case is for an assessment-sensitive folk concept of knowledge, the less reason we have to think the folk concept is the one of interest to epistemologists, then *at the same time*, we are implicitly recognizing two entirely *independent standards* – the folk standard and the epistemologist's standard. And armed with two different standards, the second-wave relativist's next move is to insist

that we have the trappings for a ‘parasitic’ new kind of epistemic relativism, parasitic in that it is premised upon the *response* aimed to demonstrating semantic relativism to be in an important respect orthogonal to mainstream epistemology. And to stress, rather than to relativise knowledge (as traditional arguments have) to different cultural epistemic frameworks or styles of reasoning (e.g. Rorty, Wittgenstein, Hacking) or (as new, semantic epistemic relativism does) to a context of assessment (e.g. MacFarlane) the ‘second-waver’ tells us that no statement of the form ‘S knows *p*’ is true, *simpliciter*, but that it receives a truth value only once a certain parameter, X, is filled in, where there are two possible values for X: the standards apposite to the *ordinary folk concept* and the *standards apposite to the epistemologists’ concept*.

28. Or perhaps merely homonymous depending on how closely related we want to grant that the two meanings are. As Murphy & Koskela (2010) note, ‘The distinction between homonymy and polysemy is usually made on the basis of the relatedness of the senses: polysemy involves related senses, whereas the senses associated with homonymous lexemes are not related.’
29. Hales (2014, 65–66), a relativist, grants this much. After all, as Hales notes, in the famous case of the dispute concerning whether a squirrel was ‘going around’ the tree – the opening of William James’ (1907, 43–45) lecture *What Pragmatism Means* – this dispute was not settled by relativising ‘going around’ to different standards, but rather, to point out that ‘going around’ is ambiguous between two distinct senses the conflation of which was responsible for the dispute.
30. Consider a parallel point. Hazlett argued, on the basis of discovering that the ordinary concept of knowledge is non-factive and the epistemologist’s concept cannot be, that these are two different things. In doing so he is not opening himself up to the following charge: that he’s *motivated* relativising knowledge to two different kinds of standards. Relativism is a profound and controversial thesis, and we don’t motivate the latter by pointing to the former which all absolutists grant as an unsurprising fact about language.
31. Kappel (2010, 71–72, with a presentational tweak). Cf. Chapter §8.2 on ‘Epistemic Aftermath’ for discussion of the value turn in epistemology.
32. My italics.
33. Thanks to Michael Hannon for directing me to these references.
34. For a systematic approach to thinking about how a norm sustains a practice, see Turri (*forthcoming*).
35. For example, in response to the envisioned ‘second-wave’ relativist’s countermove.
36. This is the *illumination thesis*, which we saw in the previous chapter that Hazlett thinks we should reject given that the ordinary concept of knowledge is, according to him, evidenced as non-factive given use data.

# References

- Adler, Jonathan. (2004) 'Reconciling Openmindedness and Belief.' *Theory and Research in Education* 2: 127–142.
- Ahlstrom-Vij, Kristoffer, and Stephen R. Grimm. (2013) 'Getting it Right.' *Philosophical Studies* 166: 329–347.
- Aikin, Scott F. (2005) 'Who Is Afraid of Epistemology's Regress Problem?' *Philosophical Studies* 126: 191–217.
- Aikin, Scott. (2011) *Epistemology and the Regress Problem*. Routledge.
- Aiken, Scott. (2014) 'Knowing Better, Cognitive Command and Epistemic Infitism.' In John Turri and Peter Klein, (eds), *Ad Infinitum: New Essays on Epistemological Infitism*. Oxford: Oxford University Press: 18–36.
- Alfano, Mark. (2012) 'Expanding the Situationist Challenge to Responsibilist Virtue Epistemology.' *The Philosophical Quarterly* 62: 223–249.
- Alfano, Mark. (2013) *Character as Moral Fiction*. Cambridge: Cambridge University Press.
- Alfano, Mark. (2014) 'Expanding the Situationist Challenge to Reliabilism About Inference.' In Abrol Fairweather (ed.), *Virtue Epistemology Naturalized*. Springer: 103–122.
- Alfano, Mark, and Don Loeb. (2014). 'Experimental Moral Philosophy.' In Edward N. Zalta, Summer (eds) *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2014/entries/experimental-moral/>.
- Alston, William. (1978) 'Meta-Ethics and Metaepistemology.' In Alvin Goldman and Jaegwon Kim (eds), *Values and Morals*. Dordrecht: D. Reidel Press: 275–297.
- Alston, William. (1993a) *Perceiving God: The Epistemology of Religious Experience*. Cornell: Cornell University Press.
- Alston, William. (1993b) *The Reliability of Sense Perception*. Cornell: Cornell University Press.
- Alston, William P. (1986) 'Epistemic Circularity.' *Philosophy and Phenomenological Research* 47: 1–30.
- Alston, William. (2005) *Beyond Justification: Dimensions of Epistemic Evaluation*. Cambridge: Cambridge University Press.
- Armstrong, David. (1973) *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Baehr, Jason. (2009) 'Evidentialism, Vice, and Virtue.' *Philosophy and Phenomenological Research* 78: 545–567.
- Baehr, Jason. (2011) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.
- Baghramian, Maria. (2004) *Relativism*. London: Routledge.
- Baghramian, Maria. (2014) *The Many Faces of Relativism*. London: Routledge.
- Baghramian, Maria, and J. Adam Carter. (2015) 'Relativism.' *Stanford Encyclopaedia of Philosophy* 1–46.
- Barnes, Jonathan. (1998) 'The Beliefs of a Pyrrhonist.' In Myles Burnyeat and Michael Frede (eds), *The Original Sceptics: A Controversy*. Indianapolis: Hackett Publishing Company: 58–91.

- Battaly, Heather. (2004) 'Must the Intellectual Virtues Be Reliable?' INPC Session Paper (manuscript).
- Beebe, James. (2012) 'Social Functions of Knowledge Attributions.' In Jessica Brown and Mikkel Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press: 220–242.
- Belnap, Nuel D., Michael Perloff, and Ming Xu. (2001) *Facing the Future: Agents and Choices in our Indeterminist World*. Oxford: Oxford University Press.
- Bengson, John, and Marc Moffett. (2011) 'Two Conceptions of Mind and Action: Knowing How and the Philosophical Theory of Intelligence.' In John Bengson and Marc Moffett (eds), *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford: Oxford University Press: 3–55.
- Berendzen, J.C. (2014) 'Motor Imagery and Merleau-Pontyan Accounts of Skilled Action.' *Ergo, an Open Access Journal of Philosophy* 1: 169–198.
- Bergmann, Michael. (1997) 'Internalism, Externalism and the No-Defeater Condition.' *Synthese* 110: 399–417.
- Bergmann, Michael. (2008) 'Externalist Responses to Skepticism.' In John Greco (ed.), *Oxford Handbook of Skepticism*. Oxford: Oxford University Press: 504–532.
- Berkeley, George. (1878) *A Treatise Concerning the Principles of Human Knowledge*. Philadelphia, PA: J.B. Lippincott & Company.
- Bertrand, Marianne, and Sendhil Mullainathan. (2003) *Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination*. National Bureau of Economic Research: Paper 9873.
- Bird, Alexander. (2013) 'Thomas Kuhn.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2013. <http://plato.stanford.edu/archives/fall2013/entries/thomas-kuhn/>.
- Blackburn, Simon. (1993) *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Bland, Steven. (2014) 'Incommensurability, Relativism, and the Epistemic Authority of Science.' *Episteme* 11: 463–473.
- Blome-Tillmann, Michael. (2009) 'Knowledge and Presuppositions.' *Mind* 118: 241–294.
- Boghossian, Paul. (1989) 'Content and Self-Knowledge.' *Philosophical Topics* 17: 5–26.
- Boghossian, Paul. (2001) 'How Are Objective Epistemic Reasons Possible?' *Philosophical Studies* 106: 1–40.
- Boghossian, Paul. (2006a) *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.
- Boghossian, Paul. (2006b) 'What Is Relativism?' In Patrick Greenough and Michael P. Lynch (eds), *Truth and Realism*. Oxford: Oxford University Press: 13–37.
- BonJour, Laurence. (1985) *The Structure of Empirical Knowledge*. Cambridge: Cambridge University Press.
- Brogaard, Berit. (2008) 'Moral Contextualism and Moral Relativism.' *The Philosophical Quarterly* 58: 385–409.
- Brown, Jessica A. (2013) 'Knowing-How: Linguistics and Cognitive Science.' *Analysis* 73: 220–227.
- Buckwalter, Wesley. (2010) 'Knowledge Isn't Closed on Saturday: A Study in Ordinary Language.' *Review of Philosophy and Psychology* 1 3: 395–406.

- Burnyeat, Myles. (1976) 'Protagoras and Self-Refutation in Plato's Theaetetus.' *The Philosophical Review* 85: 172–195.
- Carroll, Lewis. (1895) 'What the Tortoise Said to Achilles.' *Mind* 4: 278–280.
- Carter, J. Adam. (2013) 'Extended Cognition and Epistemic Luck.' *Synthese* 190: 4201–4214.
- Carter, J. Adam. (2014) 'Relativism, Knowledge and Understanding.' *Episteme* 11: 35–52.
- Carter, J. Adam, and Emma C. Gordon. (2014) 'Openmindedness and Truth.' *Canadian Journal of Philosophy* 44: 207–224.
- Carter, J. Adam, and Benjamin Jarvis. (2012) 'Against Swamping.' *Analysis* 72: 690–699.
- Carter, J. Adam, and Duncan Pritchard. (2013). 'Knowledge-How and Epistemic Luck.' *Noûs* (Online First: DOI: 10.1111/nous.12054).
- Carter, J. Adam, and Duncan Pritchard. (2014) 'Knowledge-How and Cognitive Achievement.' *Philosophy and Phenomenological Research* (Online First: DOI: 10.1111/phpr.12094).
- Carter, J. Adam, Benjamin Jarvis, and Katherine Rubin. (2013a) 'Knowledge and the Value of Cognitive Ability.' *Synthese* 190: 3715–3729.
- Carter, J. Adam, Benjamin Jarvis, and Katherine Rubin. (2013b) 'Knowledge: Value on the Cheap.' *Australasian Journal of Philosophy* 91: 249–263.
- Carter, J. Adam, Benjamin Jarvis, and Katherine Rubin. (2014) 'Varieties of Cognitive Achievement.' *Philosophical Studies* 172: 1603–1623.
- Carter, J. Adam, and Duncan Pritchard. (2014) 'Knowledge-How and Cognitive Achievement.' *Philosophy and Phenomenological Research* (Online First: DOI: 10.1111/phpr.12094).
- Carter, J. Adam, and Matthew Chrisman. (2012) 'Is Epistemic Expressivism Incompatible with Inquiry?' *Philosophical Studies* 159: 323–339.
- Carter, J. Adam, and Duncan Pritchard. (2015) 'Epistemic Situationism, Epistemic Dependence and the Epistemology of Education.' In Mark Alfano and Abrol Fairweather (eds), *Epistemic Situationism*. Oxford University Press.
- Carter, J. Adam, and Duncan Pritchard. (2015) 'Perceptual Knowledge and Relevant Alternatives.' *Philosophical Studies*, forthcoming.
- Chalmers, David J. (2011) 'Verbal Disputes.' *Philosophical Review* 120: 515–566.
- Chang, Ruth. (2002) 'The Possibility of Parity.' *Ethics* 112: 659–688.
- Chisholm, Roderick. (1973) *The Problem of the Criterion*. Milwaukee: Marquette University Press.
- Chisholm, Roderick. (1977) *Theory of Knowledge*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Chrisman, Matthew. (2007) 'From Epistemic Contextualism to Epistemic Expressivism.' *Philosophical Studies* 135: 225–254.
- Chrisman, Matthew. (2012) 'Epistemic Expressivism.' *Philosophy Compass* 7: 118–126.
- Christensen, David. (2009) 'Disagreement as Evidence: The Epistemology of Controversy.' *Philosophy Compass* 4: 756–767.
- Clifford, William Klingdon. (1879/1999) *The Ethics of Belief*. In *The Ethics of Belief and Other Essays*. New York: Prometheus Books.
- Colaço, David, Wesley Buckwalter, Stephen Stich, and Edouard Machery. (2014) 'Epistemic Intuitions in Fake-Barn Thought Experiments.' *Episteme* 11: 199–212.

- Conee, Earl, and Richard Feldman. (2004) *Evidentialism: Essays in Epistemology: Essays in Epistemology*. Oxford: Oxford University Press.
- Cuneo, Terence. (2007) *The Normative Web: An Argument for Moral Realism*. Oxford: Clarendon Press.
- Correll, Joshua, Geoffrey R Urland, and Tiffany A Ito. (2006) 'Event-Related Potentials and the Decision to Shoot: The Role of Threat Perception and Cognitive Control.' *Journal of Experimental Social Psychology* 42: 120–128.
- Crombie, Alistair. (1981) 'Philosophical Presuppositions and Shifting Interpretations of Galileo.' In *Theory Change, Ancient Axiomatics, and Galileo's Methodology*. Dordrecht: Springer. 271–286.
- Dancy, Jonathan. (1982) 'Intuitionism in Metaepistemology.' *Philosophical Studies* 42: 395–408.
- Davies, Martin. (2003) 'Armchair Knowledge, Begging the Question, and Epistemic Warrant.' *Unpublished Carl G. Hempel Lectures, Princeton University*.
- Derrida, Jacques. (1974) *Of Grammatology*. Baltimore, MD: Johns Hopkins University Press.
- DeRose, Keith. (1992) 'Contextualism and Knowledge Attributions.' *Philosophy and Phenomenological Research* 52: 913–929.
- DeRose, Keith. (2004) 'Single Scoreboard Semantics.' *Philosophical Studies* 119: 1–21.
- DeRose, Keith. (2009) *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1: Knowledge, Skepticism, and Context*. Oxford: Oxford University Press.
- deTurck, Mark A., Janet J. Harszslak, Darlene J. Bodhorn, and Lynne A. Texter. (1990) 'The Effects of Training Social Perceivers to Detect Deception from Behavioral Cues.' *Communication Quarterly* 38: 189–199.
- Devitt, Michael. (1983) 'Dummett's Anti-Realism.' *The Journal of Philosophy* 80: 73–99.
- Doris, John Michael. (2002) *Lack of Character: Personality and Moral Behavior*. Cambridge: Cambridge University Press.
- Downing, Lisa. (2013) 'George Berkeley', *The Stanford Encyclopedia of Philosophy* (Spring 2013 Edition), Edward N. Zalta (ed.), <http://plato.stanford.edu/entries/berkeley/>.
- Dreier, James. (2004) 'Meta-Ethics and the Problem of Creeping Minimalism.' *Philosophical Perspectives* 18: 23–44.
- Dretske, Fred. (1970) 'Epistemic Operators.' *The Journal of Philosophy* 67: 1007–1023.
- Dreyfus, Hubert L. (2005) 'Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise.' In *Proceedings and Addresses of the American Philosophical Association* 47–65.
- Dreyfus, Hubert L. (2007) 'The Return of the Myth of the Mental.' *Inquiry* 50: 352–365.
- Dutant, Julien. (2013) 'In Defense of Swamping.' *Thought: A Journal of Philosophy* 2: 357–366.
- Elga, Adam. (2007) 'Reflection and Disagreement.' *Noûs* 41: 478–502.
- Emilsson, Eyjólfur Kjalar. (1994) 'Plato's Self-Refutation Argument in Theaetetus 171a–c Revisited.' *Phronesis* 39: 136–149.
- Engel, Pascal. (2004) 'Truth and the Aim of Belief.' In Donald Gillies (ed.), *Laws and Models in Science. Edited by Donald Gillies*. London: King's College Publication 77–97.

- Enoch, David. (2011) *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Enoch, David, and Joshua Schechter. (2008) 'How Are Basic Belief-Forming Methods Justified?' *Philosophy and Phenomenological Research* 76: 547–579.
- Evans-Pritchard, E. E. (1937) *Witchcraft, Oracles and Magic among the Azande*. Oxford: Clarendon Press.
- Evnine, Simon. (1999) 'Believing Conjunctions.' *Synthese* 118: 201–127.
- Fantl, Jeremy. (2003) 'Modest Informatism.' *Canadian Journal of Philosophy* 33: 537–562.
- Feldman, Richard. (2000) 'The Ethics of Belief.' *Philosophical and Phenomenological Research* 60: 667–695.
- Feldman, Richard. (2007) 'Reasonable Religious Disagreements.' In Louise Anthony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular*. Oxford: Oxford University Press: 194–214.
- Field, Hartry. (1998) 'Epistemological Nonfactualism and the A Prioricity of Logic.' *Philosophical Studies* 92: 1–24.
- Fine, Kit. (2001) 'The Question of Realism.' *Philosophers' Imprint* 1:1–30.
- Frede, Michael. (1998a) 'The Sceptic's Beliefs.' In Myles Burnyeat and Michael Frede (eds), *The Original Sceptics: A Controversy*. Indianapolis: Hackett Publishing Company.
- Frede, Michael. (1998b) 'The Sceptic's Two Kinds of Assent and the Question of the Possibility of Knowledge.' In Myles Burnyeat and Michael Frede (eds), *The Original Sceptics: A Controversy*. Indianapolis: Hackett Publishing Company.
- Frege, Gottlob. (1979) 'Posthumous Writings' in H. Hermes, F. Kambartel and F. Kaulbach (eds), *Trans. P. Long and R. White*. Oxford: Basil Blackwell.
- Fumerton, Richard A. (1995) *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield Publishers.
- Gelfert, Axel. (2011) 'Expertise, Argumentation, and the End of Inquiry.' *Argumentation* 25: 297–312.
- Geach, Peter T. (1965) 'Assertion.' *The Philosophical Review* 74: 449–465.
- Gerken, Mikkel. (2012a) 'Discursive Justification and Skepticism.' *Synthese* 189: 373–394.
- Gerken, Mikkel. (2012b) 'Univocal Reasoning and Inferential Presuppositions.' *Erkenntnis* 76: 373–394.
- Gerken, Mikkel. (2015) 'How to Do Things with Knowledge Ascriptions.' *Philosophy and Phenomenological Research* 90: 223–234.
- Gibbard, Allan. (1992) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Glick, Ephraim. (2011) 'Two Methodologies for Evaluating Intellectualism.' *Philosophy and Phenomenological Research* 83: 398–434.
- Goldberg, Sanford C. (2015) 'Should have Known.' *Synthese*: 1–32. (Online First: DOI: 10.1007/s11229-015-0662-z).
- Goldman, Alvin. (1967) 'A Causal Theory of Knowing.' *The Journal of Philosophy* 64: 357–372.
- Goldman, Alvin. (1976) 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy* 73: 771–791.
- Goldman, Alvin I. (1979) 'What is Justified Belief?' In Ernest Sosa & Jaegwon Kim (eds), *Epistemology. An Anthology*. London: Blackwell. 340–353.
- Goldman, Alvin. (1999) *Knowledge in a Social World*. Oxford: Oxford University Press.

- Goldman, Alvin I. (2008) 'Immediate Justification and Process Reliabilism'. In Quentin Smith (ed.), *Epistemology: New Essays*. Oxford: Oxford University Press. 63–82.
- Goldman, Alvin. (2011) 'Reliabilism.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2011. <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Goldman, Alvin., and E. Olsson. (2009) 'Reliabilism and the Value of Knowledge.' In Adrian Haddock, Alan Millar and Duncan Pritchard (eds), *Epistemic Value*. Oxford: Oxford University Press.
- Goodale, Melvyn A, and A David Milner. (1992) 'Separate Visual Pathways for Perception and Action.' *Trends in Neurosciences* 15: 20–25.
- Goodman, Nelson. (1975) 'Words, Works, Worlds.' *Erkenntnis* 9: 57–73.
- Goodman, Nelson. (1978) *Ways of Worldmaking*. Indianapolis, IN: Hackett Publishing Company.
- Greco, John. (2010) *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Greco, John. (2013) 'Reflective Knowledge and the Pyrrhonian Problematic.' In John Greco and John Turri (eds), *Virtuous Thoughts: The Philosophy of Ernest Sosa*. Dordrecht: Springer: 179–191.
- Greenwald, Anthony G., Mark A. Oakes, and Hunter G. Hoffman. (2003) 'Targets of Discrimination: Effects of Race on Responses to Weapons Holders.' *Journal of Experimental Social Psychology* 39: 399–405.
- Grice, Paul. (1991) *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Grimm, Stephen R. (2014) 'Understanding as Knowledge of Causes.' In Abrol Fairweather (ed), *Virtue Epistemology Naturalized*. Dordrecht: Springer: 329–345.
- Grote, George. (1865) *Plato, and the Other Companions of Sokrates*. Vol. 2. J. Murray.
- Hacking, Ian. (1982) 'Language, Truth and Reason.' In *Rationality and Relativism*, 48–66.
- Hage, Jaap C. (1997) *Reasoning with Rules*. Dordrecht: Springer.
- Hales, Steven D. (1997) 'A Consistent Relativism.' *Mind* 106: 33–52.
- Hales, Steven D. (2006) *Relativism and the Foundations of Philosophy*. Cambridge, MA: MIT Press.
- Hales, Steven D. (2014) 'Motivations for Relativism as a Solution to Disagreements.' *Philosophy* 89: 63–82.
- Harman, Gilbert. (1975) 'Moral Relativism Defended.' *Philosophical Review* 84: 3–22.
- Harman, Gilbert. (1996) *Moral Relativism and Moral Objectivity*. Oxford: Blackwell Publishing.
- Harman, Gilbert. (1999) 'Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error.' *Proceedings of the Aristotelian Society* 99: 315–331.
- Harman, Gilbert. (2000) 'The Nonexistence of Character Traits.' *Proceedings of the Aristotelian Society* 100:223–226.
- Hájek, Alan. (2014) 'Philosophical Heuristics and Philosophical Creativity.' *The Philosophy of Creativity: New Essays*. Oxford: Oxford University Press.

- Hazlett, Allan. (2010) 'The Myth of Factive Verbs.' *Philosophy and Phenomenological Research* 80: 497–522.
- Hills, Alison. (2009) 'Moral Testimony and Moral Epistemology.' *Ethics* 120: 94–127.
- Huemer, Michael. (2005) *Ethical Intuitionism*. Palgrave Macmillan.
- Ichikawa, Jonathan Jenkins. (2013) 'Basic Knowledge and Contextualist "E=K".' *Thought: A Journal of Philosophy* 2: 282–292.
- Ichikawa, Jonathan Jenkins and Steup, Matthias. 'The Analysis of Knowledge' *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>>.
- Jenkins, Carrie (2005) 'Realism and Independence.' *American Philosophical Quarterly* 42: 199–209.
- Jenkins, Carrie Ichikawa (2014) 'Merely Verbal Disputes.' *Erkenntnis* 79: 11–30.
- John MacFarlane (2011) Relativism and Knowledge Attributions. In Duncan Pritchard & Sven Bernecker (eds), *The Routledge Companion to Epistemology*. Routledge. 536–544.
- Jordan, James. (1971) 'Protagoras and Relativism: Criticisms Good and Bad.' *Southwestern Journal of Philosophy* 2: 1–27.
- Joyce, Richard. (2007) *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kahane, Guy. (2011) 'Evolutionary Debunking Arguments.' *Noûs* 45: 103–125.
- Kallestrup, Jesper. (2012) 'Bootstrap and Rollback: Generalizing Epistemic Circularity.' *Synthese* 189: 395–413.
- Kamp, Hans. (1971) 'To the Memory of Arthur Prior: Formal Properties of 'Now'.' *Theoria* 37: 227–273.
- Kaplan, David. (1989) 'Demonstratives'. In J. Almog, J. Perry, and H. Wettstein (eds), *Themes from Kaplan*. Oxford: Oxford University Press: 481–563.
- Kappel, Klemens. (2010) 'On Saying that Someone Knows.' In Adrian Haddock, Alan Millar and Duncan Pritchard (eds), *The Nature and Value of Knowledge*. Oxford: Oxford University Press, Oxford, 69–88.
- Keller, John A. (2014) 'Theological Anti-Realism.' *Journal of Analytic Theology* 2: 13–42.
- Kelly, Thomas. (2005) 'The Epistemic Significance of Disagreement.' *Oxford Studies in Epistemology* 1: 167–196.
- Kelp, Christoph. (2011) 'What's the Point of "Knowledge" Anyway?' *Episteme* 8: 53–66.
- Kelp, Christoph. (2014) 'Two for the Knowledge Goal of Inquiry.' *American Philosophical Quarterly* 51: 227–232.
- Klein, Peter. (2003) 'When Infinite Regresses Are Not Vicious.' *Philosophy and Phenomenological Research* 66: 718–729.
- Klein, Peter. (2007) 'How to Be an Infinitist About Doxastic Justification.' *Philosophical Studies* 134: 25–29.
- Klein, Peter D. (1999) 'Human Knowledge and the Infinite Regress of Reasons.' *Noûs* 33: 297–325.
- Knobe, Joshua, and Shaun Nichols. (2008) *Experimental Philosophy*. Oxford: Oxford University Press.
- Kölbel, Max. (2004) 'Faultless Disagreement.' *Proceedings of the Aristotelian Society* 104:53–73.

- Kölbel, Max. (2011) 'Global Relativism and Self-Refutation.' In Steven Hales (ed), *A Companion to Relativism*. Oxford: Blackwell Publishing: 11–30.
- Kuhn, Thomas. (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kupreeva, Inna. (2014) 'Socrates' Refutation of Protagoras' 1–30 (manuscript).
- Kusch, Martin. (2009) 'Boghossian on Epistemological and Moral Relativism: A Critique.' *Unpublished Manuscript*.
- Kusch, Martin. (2010) 'Epistemic Replacement Relativism Defended.' In Mauricio Suarez (ed.), *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*: 165–175.
- Kvanvig, Jonathan L. (2013) 'Curiosity and a Response-Dependent Account of the Value of Understanding,' In Tim Henning and David Schweikard (eds), *Knowledge, Virtue, and Action*. Boston: Routledge Publishing Co.: 151–175.
- Lackey, Jennifer. (1999) 'Testimonial Knowledge and Transmission.' *Philosophical Quarterly* 50: 71–490.
- Lackey, Jennifer. (2005) 'Memory as a Generative Epistemic Source.' *Philosophy and Phenomenological Research* 70: 636–658.
- Lackey, Jennifer. (2008) *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lackey, Jennifer. (2010) 'Testimonial Knowledge.' In Sven Bernecker & Duncan Pritchard (eds), *Routledge Companion to Epistemology*, London: Routledge: 316–325.
- Lammenranta, Markus. (2008) 'The Pyrrhonian Problematic.' In John Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford University Press: 9–33.
- Latour, Bruno and Steve Woolgar. (1986) *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press.
- Lawlor, Krista. (2013) *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.
- Lee, Mi-Kyoung. (2005) *Epistemology After Protagoras: Responses to Relativism in Plato, Aristotle, and Democritus*. Oxford: Oxford University Press.
- Lehrer, Keith. (1971) 'Why Not Scepticism?' *The Philosophical Forum* 2:283–298.
- Lewis, David. (1980) 'Index, Context, and Content.' In Stig Kanger & Sven Öhman (eds), *Philosophy and Grammar*. Reidel. 79–100.
- Lipton, Peter. (2004) *Inference to the Best Explanation*. 2nd edn New York: Routledge.
- Littlejohn, Clayton. (2013) 'Disagreement and Defeat.' In Diego Machuca (ed.), *Disagreement and Skepticism*. London: Routledge: 169–193.
- Lynch, Michael. (2010) 'Epistemic Circularity and Epistemic Incommensurability.' In Adrian Haddock, Alan Millar, and Duncan Pritchard (eds), *Social Epistemology*. Oxford: OUP: 262–277.
- Maccallum, Gerald C. (1966) 'On Applying Rules.' *Theoria* 32: 196–210.
- MacFarlane, John. (2005a) 'The Assessment Sensitivity of Knowledge Attributions.' *Oxford Studies in Epistemology* 1: 197–233.
- MacFarlane, John. (2005b) 'Making Sense of Relative Truth.' *Proceedings of the Aristotelian Society* 105:305–323.
- MacFarlane, John. (2007) 'Relativism and Disagreement.' *Philosophical Studies* 132: 17–31.
- MacFarlane, John. (2014) *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

- Madison, Brent. (2011) 'Combating Anti Anti-luck Epistemology.' *Australasian Journal of Philosophy* 89: 47–58.
- Maffie, James. (1993) 'Realism, Relativism, and Naturalized Metaepistemology.' *Metaphilosophy* 24: 1–13.
- Markie, Peter J. (2005) 'Easy Knowledge.' *Philosophy and Phenomenological Research* 70: 406–416.
- Matthen, Mohan. (1985) 'Perception, Relativism, and Truth: Reflections on Plato's *Theaetetus* 152–160.' *Dialogue* 24: 33–58.
- McCain, Kevin. (2013) 'Explanationist Evidentialism.' *Episteme* 10: 299–315.
- McCain, Kevin. (2014a) *Evidentialism and Epistemic Justification*. London: Routledge.
- McCain, Kevin. (2014b) 'Problem of the Criterion.' *Internet Encyclopedia of Philosophy*.
- McGlynn, Aidan. (2013) 'Believing Things Unknown.' *Noûs* 47: 385–407.
- McKenna, Robin. (2014) 'Shifting Targets and Disagreements.' *Australasian Journal of Philosophy* 92: 725–742.
- McKenna, Robin. (2015a) 'Epistemic Contextualism Defended.' *Synthese* 192: 363–383.
- McKenna, Robin. (2015b) 'Contextualism in Epistemology.' *Analysis Reviews*, forthcoming.
- Meiland, Jack. (1979) 'Is Protagorean Relativism Self-Refuting?' *Grazer Philosophische Studien* 9: 51–68.
- Meiland, Jack and Michael Krausz. (1982) (eds) *Relativism, Cognitive and Moral*. Notre Dame, Indiana: University of Notre Dame Press.
- Miller, Alexander. (2012) 'Realism.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2012. <http://plato.stanford.edu/archives/spr2012/entries/realism/>.
- Moore, George Edward. (1925) *A Defence of Common Sense*. Sydney, Australia: Allen & Unwin.
- Moore, George Edward. (1939) 'Proof of an External World.' *Proceedings of the British Academy* 25: 273–300.
- Moore, George Edward, and Casimir Lewy. (1962) *Commonplace Book 1919–1953*.
- Moss-Racusin, Corinne A, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman. (2012) 'Science Faculty's Subtle Gender Biases Favor Male Students.' *Proceedings of the National Academy of Sciences* 109: 16474–16479.
- Mounce, H. O. (2007) 'Wittgenstein and Classical Realism.' *Readings of Wittgenstein's On Certainty*, ed. Benner, W.H. and Moyal-Sharrock, D. Palgrave, pp. 103–121.
- Murphy, M. Lynne and Anu Koskela. (2010) *Key Terms in Semantics*. New York; London: Continuum International Publishing.
- Nagel, Thomas. (2001) *The Last Word*. New York: Oxford University Press.
- Neta, Ram (2007) 'Fixing the Transmission: The New Mooreans.' In Susana Nuccetelli & Gary Seay (eds), *Themes From G. E. Moore: New Essays in Epistemology and Ethics*. Oxford: Clarendon Press: 62–83.
- Nozick, Robert. (2001) *Invariances: The Structure of the Objective World*. Cambridge, MA: Harvard University Press.
- Oberheim, Eric, and Paul Hoyningen-Huene. (2013) 'The Incommensurability of Scientific Theories.' In *The Stanford Encyclopedia of Philosophy*, edited by

- Edward N. Zalta, Spring 2013. <http://plato.stanford.edu/archives/spr2013/entries/incommensurability/>.
- O'Grady, Paul. (2002) *Relativism*. Montreal: McGill-Queen's Press.
- Olson, Jonas (2011) 'Error Theory and Reasons for Belief.' In Andrew Reisner & Asbjørn Steglich-Petersen (eds), *Reasons for Belief*. Cambridge University Press.
- Olsson, Erik J. (2007) 'Reliabilism, Stability, and the Value of Knowledge.' *American Philosophical Quarterly* 44: 343–355.
- Olsson, Erik. (2014) 'Coherentist Theories of Epistemic Justification.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2014. <http://plato.stanford.edu/archives/spr2014/entries/justep-coherence/>.
- Parfit, Derek. (2013) *On What Matters*. Oxford: Oxford University Press.
- Parsons, Josh. (2013) 'Presupposition, Disagreement, and Predicates of Taste.' *Proceedings of the Aristotelian Society* 113:163–173.
- Passmore, John. (1970) *Philosophical Reasoning*. London: Duckworth.
- Peacocke, Christopher. (1999) *Being Known*. Oxford, Oxford University Press.
- Peters, Douglas P, and Stephen J Ceci. (1982) 'Peer-Review Practices of Psychological Journals: The Fate of Published Articles, Submitted Again.' *Behavioral and Brain Sciences* 5:187–195.
- Pollock, John. (1986) *Contemporary Theories of Knowledge*. Savage, MD: Rowman and Littlefield.
- Pritchard, Duncan. (2002) 'Recent Work on Radical Skepticism.' *American Philosophical Quarterly* 39: 215–257.
- Pritchard, Duncan. (2004) 'Some Recent Work in Epistemology.' *The Philosophical Quarterly* 54: 604–13.
- Pritchard, Duncan. (2008) 'On Metaepistemology.' *Identity* 16: 645–672.
- Pritchard, Duncan. (2010a) 'Epistemic Relativism, Epistemic Incommensurability, and Wittgensteinian Epistemology.' In Steven Hales (ed.), *A Companion to Relativism*. Wiley Blackwell: 266–285.
- Pritchard, Duncan. (2010b) 'Relevant Alternatives, Perceptual Knowledge and Discrimination.' *Noûs* 44: 245–268.
- Pritchard, Duncan. (2012) 'Wittgenstein and the Groundlessness of our Believing.' *Synthese* 189: 255– 272. 10.1007/s11229-011-0057-8.
- Pritchard, D., Millar, A., and Haddock, A. (2010) *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- Pryor, James. (2000) 'The Skeptic and the Dogmatist.' *Noûs* 34: 517–549.
- Pryor, James. (2004) 'What's Wrong with Moore's Argument?' *Philosophical Issues* 14: 349–378.
- Riggs, Wayne. (2008) 'The Value Turn in Epistemology.' In Vincent Hendricks (ed.), *New Waves in Epistemology*. Palgrave Macmillan. 300–323.
- Riggs, Wayne. (2010) 'Open-Mindedness.' *Metaphilosophy* 41: 172–188.
- Roberts, Robert, & Wood, Jay. (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Clarendon Press.
- Rorty, Richard. (1979) *Philosophy and the Mirror of Nature*. Princeton, N.J.: Princeton University Press.
- Rothschild, Daniel. (2007) 'Presuppositions and Scope.' *The Journal of Philosophy* 104: 71–106.
- Ryle, Gilbert. (1945) 'Knowing How and Knowing That: The Presidential Address.' In *Proceedings of the Aristotelian Society* 46: 1–16.

- Rysiew, Patrick. (2011). 'Epistemic Contextualism' *The Stanford Encyclopedia of Philosophy* (Winter 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>>.
- Salmon, Wesley. (1989) 'Four Decades of Scientific Explanation.' In Philip Kitcher and Wesley Salmon (eds), *Minnesota Studies in the Philosophy of Science*, vol. 13. Minneapolis: University of Minnesota Press.
- Sankey, Howard. (1999) 'Incommensurability – An Overview.' *Divinatio: Studia*.
- Sankey, Howard. (2010) 'Witchcraft, Relativism and the Problem of the Criterion.' *Erkenntnis* 72: 1–16.
- Sankey, Howard. (2011) 'Epistemic Relativism and the Problem of the Criterion.' *Studies in History and Philosophy of Science Part A* 42: 562–570.
- Sankey, Howard. (2012) 'Scepticism, Relativism and the Argument from the Criterion.' *Studies in History and Philosophy of Science Part A* 43: 182–190.
- Saul, Jennifer. (2013) 'Scepticism and Implicit Bias.' *Disputatio* 5: 243–263.
- Sayre-McCord, Geoff. (2014) 'Metaethics.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2014. <http://plato.stanford.edu/archives/sum2014/entries/metaethics/>.
- Sayre-McCord, Geoff. (1991). 'Being a Realist About Relativism (in Ethics).' *Philosophical Studies* 61: 155–176.
- Schmitt, Frederick. (2004) 'What Is Wrong with Epistemic Circularity?' *Philosophical Issues* 14: 379–402.
- Schaffer, Jonathan. (2001). 'Knowledge, Relevant Alternatives and Missed Clues.' *Analysis* 61: 202–208.
- Schaffer, Jonathan. (2005) 'Contrastive Knowledge.' In T. Gendler & J. Hawthorne (eds), *Oxford Studies in Epistemology*. Oxford: Oxford University Press.
- Schaffer, Jonathan. (2014) 'Monism.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2014. <http://plato.stanford.edu/archives/spr2014/entries/monism/>.
- Seidel, Markus. (2013) 'Why the Epistemic Relativist Cannot Use the Sceptic's Strategy. A Comment on Sankey.' *Studies in History and Philosophy of Science Part A* 44: 134–139.
- Shafer-Landau, Russ. (2005) *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Shah, Nishi. (2008) 'Denying Doxastic Reasons.' [http://normativityofbeliefworkshop.weebly.com/uploads/2/8/1/2/2812762/shah\\_paper.pdf](http://normativityofbeliefworkshop.weebly.com/uploads/2/8/1/2/2812762/shah_paper.pdf)
- Sider, Theodore. (2002) 'The Ersatz Pluriverse.' *The Journal of Philosophy* 99: 279–315.
- Siegel, Harvey. (2011) 'Epistemological Relativism: Arguments Pro and Con.' In Steven Hales (ed.), *A Companion to Relativism*. Wiley Blackwell: 199–218.
- Slovic, Paul. (1987) 'Perception of Risk.' *Science* 236: 280–285.
- Sober, Elliott. (1982) 'Realism and Independence.' *Noûs* 16: 369–385.
- Sosa, Ernest. (1991) *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Sosa, Ernest. (1999) 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* 13: 141–154.
- Sosa, Ernest. (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.
- Sosa, Ernest. (2009) *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*. Oxford: Oxford University Press.

- Spencer, Jack. (2014) 'Relativity and Degrees of Relationality.' *Philosophy and Phenomenological Research*.
- Srinivasan, Amia. (2014) 'The Archimedean Urge.' Unpublished draft.
- Srinivasan, Amia. (2014) 'What's in a Norm?' Unpublished draft.
- Stalnaker, Robert. (1973) 'Presuppositions.' *Journal of Philosophical Logic* 2: 447–457.
- Stalnaker, Robert. (1974) 'Pragmatic Presuppositions.' In Robert Stalnaker (ed), *Context and Content*. Oxford: Oxford University Press.
- Robert Stalnaker (1978) 'Assertion.' *Syntax and Semantics* 9: 315–332.
- Stalnaker, Robert. (2002) 'Common Ground.' *Linguistics and Philosophy* 25: 701–721.
- Stalnaker, Robert. (2008) 'A Response to Abbott on Presupposition and Common Ground.' *Linguistics and Philosophy* 31: 539–544.
- Stanley, Jason. (2011) *Know How*. Oxford: Oxford University Press.
- Stanley, Jason, and Timothy Williamson. (2001) 'Knowing How.' *The Journal of Philosophy* 98: 411–444.
- Stanley, Jason, and John W Krakauer. (2013) 'Motor Skill Depends on Knowledge of Facts.' *Frontiers in Human Neuroscience* 7.
- Strawson, Peter F. (1985) *Skepticism and Naturalism: Some Varieties*. Cambridge: Cambridge University Press.
- Streumer, Bart. (2012) 'Can We Believe the Error Theory?' *Journal of Philosophy* 110: 1–25.
- Strevens, Michael. (2008) *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Stroud, Barry. (1984) *The Significance of Philosophical Scepticism*. Oxford University Press.
- Sudduth, Michael. (200). 'Defeaters in Epistemology.' *Internet Encyclopedia of Philosophy*.
- Swoyer, Chris. (2014) 'Relativism.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2014. <http://plato.stanford.edu/archives/sum2014/entries/relativism/>.
- Toribio, Josefa. (2008) 'How Do We Know How?' *Philosophical Explorations* 11: 39–52.
- Turri, John. (2010) 'Epistemic Invariantism and Speech Act Contextualism.' *Philosophical Review* 119: 77–95.
- Turri, John. (2011a) 'Manifest Failure: The Gettier Problem Solved.' *Philosopher's Imprint* 11: 1–11.
- Turri, John. (2011b) 'Mythology of the Factive.' *Logos & Episteme* 2: 143–152.
- Turri, John, and Peter D Klein. (2014) *Ad Infinitum: New Essays on Epistemological Informatism*. Oxford: Oxford University Press.
- Turri, John, Wesley Buckwalter, and Peter Blouw. (2014) 'Knowledge and Luck.' *Psychonomic Bulletin & Review* 1–13.
- Turri, John. (2015) 'Epistemic Situationism and Cognitive Ability.' In Mark Alfano and Abrol Fairweather (eds), *Epistemic Situationism*. Oxford: Oxford University Press.
- Turri, John. Sustaining Rules: A Model and Application. In J. A. Carter, E. C. Gordon, & B. Jarvis (eds), *Knowledge first*. Oxford University Press. (Forthcoming).

- Tversky, Amos, and Daniel Kahneman. (1973) 'Availability: A Heuristic for Judging Frequency and Probability.' *Cognitive Psychology* 5: 207–232.
- Tversky, Amos, and Daniel Kahneman. (1974) 'Judgment Under Uncertainty: Heuristics and Biases.' *Science* 185: 1124–1131.
- Unger, Peter. (1971) 'A Defense of Skepticism.' *The Philosophical Review* 80: 198–219.
- Unger, Peter. (2012) 'Empty Ideas.' *The Philosophers' Magazine* 57: 31–41.
- Unkelbach, Christian, Joseph P. Forgas, and Thomas F Denson. (2008) 'The Turban Effect: The Influence of Muslim Headgear and Induced Affect on Aggressive Responses in the Shooter Bias Paradigm.' *Journal of Experimental Social Psychology* 44: 1409–1413.
- Virvidakis, Stelios. (2008) 'Varieties of Quietism.' *Philosophical Inquiry* 30: 157–175.
- Vlastos, Gregory. (1956) 'Introduction to Plato's Protagoras.' *Library of Liberal Arts (New York 1956) at Xxxviii–Xlv*.
- Vogel, Jonathan. (2000) 'Reliabilism Leveled.' *The Journal of Philosophy* 97: 602–623.
- Vogel, Jonathan (2005) 'The Refutation of Skepticism'. In Steup Matthias & Sosa Ernest (eds), *Contemporary Debates in Epistemology*. Blackwell. 72–84.
- Vogt, Katja. (2014) 'Ancient Skepticism.' In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2014. <http://plato.stanford.edu/archives/win2014/entries/skepticism-ancient/>.
- William G. Lycan (2006) 'On the Gettier Problem Problem.' In Stephen Cade Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press. 148–168.
- Williams, Michael. (1991) *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Princeton, NJ: Princeton University Press.
- Williams, Michael. (2007) 'Why (Wittgensteinian) Contextualism Is Not Relativism.' *Episteme* 4: 93–114.
- Williamson, Timothy. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wittgenstein, Ludwig. (1969) *On Certainty*. Vol. 174. Oxford: Blackwell Publishing.
- Wong, David B. (1984) *Moral Relativity*. Oakland, CA: University of California Press.
- Wong, David B. (2006) *Natural Moralities: A Defense of Pluralistic Relativism*. Oxford: Oxford University Press.
- Woodward, James. (2003) *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- Wright, Crispin. (1992) *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- Wright, Crispin. (2006) 'What Is Relativism?' In Patrick Greenough and Michael P. Lynch (eds), *Truth and Realism*. New York: Oxford University Press: 13–37.
- Wright, Crispin. (2007a) 'The Perils of Dogmatism.' In Nuccetelli and Seay (eds), *Themes from GE Moore: New Essays in Epistemology and Ethics*. Clarendon Press Oxford, 25–48.

- Wright, Crispin. (2007b) 'New Age Relativism and Epistemic Possibility: The Question of Evidence.' *Philosophical Issues* 17: 262–283.
- Wright, Crispin. (2008a) 'Relativism About Truth Itself: Haphazard Thoughts About the Very Idea.' In Manuel García-Carpintero and Max Kölbel (eds), *Relative Truth*. Oxford: Oxford University Press: 157–186.
- Wright, Crispin. (2008b) 'Fear of Relativism?' *Philosophical Studies* 141: 379–390.
- Wright, Stephen. (2013) 'Does Klein's Informatism Offer a Response to Agrippa's Trilemma?' *Synthese* 190: 1113–1130.
- Zagzebski, Linda. (1996) *Virtues of the Mind*. Cambridge: Cambridge University Press.

# Index

- achievement, cognitive, 200, 255, 271–2
- action, norms of, 202, 204–5
- Adler, Jonathan, 271–2
- Aenesidemus, 246
- agnosticism, 69, 74, 101–3, 245
- Agrippa's trilemma, 62, 65, 73–5, 77–8, 247, 250, 253–4
- Ahlstrom-Vij, Kristoffer, 272
- Aikin, Scott F., 246
- Alfano, Mark, 92–3, 251
- Alston, William, 2–7, 125, 130, 233–4, 255, 257, 260
- anti-individualism, epistemic, 164, 191, 193, 269
- anti-realism
- metaepistemological, 3, 21, 23–5, 27, 217, 222, 274, 276
  - metaphysical, 17, 36, 43–7, 49–51, 53–4, 57–8, 238, 244
- Aristotle, 61, 81–2, 198
- arity, 34–5, 37, 52–3, 57, 166
- Armstrong, David, 233
- assertion
- contents of, 167, 239
  - and dialectical self-refutation, *see* self-refutation
  - self-refutation
- dogmatic, 119–20, 261
- and knowledge ascription, *see* knowledge, ascriptions of
  - norms of, 202, 205, 233, 258, 263
  - and pragmatic presupposition, *see* pragmatic presupposition
- Azande, the, 72, 74–5, 120–2, 256, 258
- Baehr, Jason, 201, 268, 271
- Baghramian, Maria, 33, 240–1, 250, 263–4
- Barnes, Jonathan, 74–5, 245, 250
- Battaly, Heather, 271
- Beebe, James, 228
- belief
- arationality of, *see* hinge propositions
  - basic, *see* basic knowledge
  - and cognitivism, 3, 239
  - and common ground, 236
  - dependence on, *see* mind-independence
  - historical contingency of, 111–12, 254
  - iteration, 10, 56
  - and knowledge, propositional, *see* knowledge, analysis of
  - and luck, epistemic, *see* knowledge, and epistemic luck
  - norms of, 201–5, 258
  - and paradox, *see* self-refutation
  - and Pyrrhonism, 60, 69, 73, 78, 245, 249, 254
  - reasons for, 23, 55, 60, 130, 219, 220, 246, 275, *see also* reasons, epistemic
  - responsible, 268
- Bellarmino, Cardinal Robert, 74, 77, 79, 81–4, 86–90, 95, 98–9, 103–4, 110, 141, 194, 251–2
- Bengson, John, and Marc Moffett, 251, 270
- Berendzen, J.C., 251
- Bergmann, Michael, 257, 268
- Berkeley, George, *see* idealism, Berkeleyan
- bias, cognitive
- implicit, 90, 93–7, 99, 252–3
  - situational and environmental, *see* situationism
- Bird, Alexander, 257
- Blackburn, Simon, 16–17
- Bland, Steven, 257
- Blome-Tillmann, Michael, 236, 255, 257
- Boghossian, Paul, 31, 34–5, 37, 41, 46, 49, 52–5, 57, 76, 79, 118, 125, 128–30, 132, 137–51, 153–4, 156–8, 161, 165, 168–9, 199–200, 240, 253–5, 258–64, 271, 274

- BonJour, Laurence, 246  
 bootstrapping, 117–18, 120–1, 257–9  
 Brogaard, Berit, 263  
 Brown, Jessica, 251, 264–5  
 Buckwalter, Wesley, 265, 267  
 Burnyeat, Myles, 31, 36–7, 39–54, 57–8, 137–9, 155, 157, 241–4, 260
- Carroll, Lewis, 130, 260  
 Carter, J. Adam, 176, 237, 240, 250–2, 256–7, 260, 263–5, 267–72  
 Chalmers, David J., 235–6, 240  
 Chang, Ruth, 270  
 Chisholm, Roderick, 6, 234, 246–7, 266  
 Chrisman, Matthew, 233, 239–40, 262, 275  
 Christensen, David., 253, 275  
 circularity, epistemic, 62, 65, 73, 78, 106–11, 113–25, 127–9, 131–6, 194, 212, 233, 248, 253–4, 256–60, 263–70  
 Clifford, William Klingdon, 185, 193, 268  
 coherentism, epistemological  
   holistic, 66, 246, 249  
   linear, 66, 246  
 Coliva, Annalisa, 256  
 common ground, 1–3, 10–16, 19–21, 110, 152, 208–10, 214, 236, 238, 273  
 Conee, Earl, and Feldman, Richard, 233, *see also* evidentialism  
 constructivism, 103, 215, 274  
 contextualism  
   attributor, 166–72, 174, 193, 211, 246, 260, 263–6, 269, 276  
   moral, 34, 149, 165, 264  
   Williams-style, 238, 246  
 criterion, problem of the, 61–2, 67–8, 246–8  
 Crombie, Alistair, 256, 277  
 Cuneo, Terence, 22–3, 26, 29, 217–20, 235–40, 275, 277
- Dancy, Jonathan, 2, 5–7, 234  
 Davies, Martin, 126  
 defeaters, epistemic  
   normative, 184–5, 188, 193, 255, 268  
   psychological, 111–12, 184, 255, 268
- DeRose, Keith, 14, 20, 264–5, 269, 272–3  
 Derrida, Jacques, 32  
 Devitt, Michael, 238, 269  
 disagreement  
   faultless, 105  
   and metaepistemological  
   commitment, 2–3, 7, 15, 208–9, 222, 234–5  
   nature of, 8, 16, 235–6, 240–1, 251, 253  
   and neutrality, 81–8, 105, 255  
   non-actual, 90, 97–100, 253  
   peer, 101–3  
   and presupposition, 8–11, 14, 15, 237–8  
   and Pyrrhonian scepticism, 73–6, 78–80, 250  
   and relativism about knowledge ascriptions, *see* knowledge, ascriptions of triggered, 95–7  
 Doris, John Michael, 91–2, 251  
 Downing, Lisa, 274–5  
 Dreier, James, 165, 238, 275  
 Dretske, Fred, 112, 275  
 Dreyfus, Hubert L., 86, 251  
 Dummett, Michael, 215, 238, 274  
 Dutant, Julien, 272
- Elga, Adam, 102, 253  
 Emilsson, Eyjólfur Kjalar, 242  
 Engel, Pascal, 272  
 Enoch, David, 233, 259  
 entitlement, epistemic, 134–5, 255, 260  
 epistemic anti-individualism, *see* anti-individualism, epistemic  
 epistemic luck, *see* knowledge, and epistemic luck  
 error-theory  
   epistemic, 21, 22, 23–7, 215–16, 218, 220, 239–40, 275  
   moral, 23, 238  
 Evans-Pritchard, E. E., 120, 258  
 evidentialism, epistemic, 5, 234, 268  
 Evnine, Simon, 243  
 expressivism, 4, 239, 241, 264  
 externalism, epistemic, 7, 61, 116–19, 122, 152, 247, 254, 257, 270

- Fantl, Jeremy, 65  
 Feldman, Richard, 5, 7–8, 14, 102,  
 152, 201, 219, 223, 233–7, 253,  
 274–6  
 Field, Hartry, 233, 239  
 Fine, Kit, 238  
 foundationalism, epistemological, 65,  
 66, 70–1, 246, 249, 256  
 Frede, Michael, 59, 245–6, 250  
 Frege, Gottlob, 8–9, 233, 238  
 Fumerton, Richard A., 122, 235,  
 258, 271  
 Galileo, 74, 77, 79, 81–4, 87–90, 95, 99,  
 103–4, 110, 139–45, 194, 220, 251  
 Geach, Peter T., 233, 249  
 Gelfert, Axel, 272  
 Gerken, Mikkel, 120, 228–9, 236–7  
 Gettier cases, *see* knowledge, and  
 epistemic luck  
 Gibbard, Allan, 16, 233, 238–9  
 Glick, Ephraim, 251  
 Goldberg, Sanford C., 185, 268  
 Goldman, Alvin, 4–8, 14, 152, 209,  
 219, 223, 233–7, 267, 272, 274,  
 276  
 Goodman, Nelson, 215, 274  
 Gordon, Emma C., 256–7, 271–2  
 Greco, John, 189, 200–1, 233, 237,  
 245–6, 269, 271  
 Grice, Paul, 3, 8, 10, 15, 236, 242, 273  
 Grimm, Stephen R., 198, 227, 271–2  
 Grote, George, 41, 44  
 ground, common, *see* common ground  
 Hacking, Ian, 108, 115, 135, 256–7,  
 263, 278  
 Hage, Jaap C., 258  
 Hájek, Alan, 49  
 Hales, Steven D., 48, 76, 80, 87–90,  
 95–101, 104–5, 108, 241, 244,  
 247, 250, 252  
 Harman, Gilbert, 34, 91, 149, 159,  
 165, 251, 264, 278  
 Hazlett, Allan, 192, 196, 206–10, 213,  
 221, 225–6, 270, 273–4, 277, 278  
 Hills, Alison, 271  
 hinge propositions, 54–5, 61, 115–16,  
 245, 256, 276–7  
 Ichikawa, Jonathan Jenkins, 199, 233,  
 235, 270–1  
 idealism, Berkeleyan, 24–7, 215–19,  
 240, 275–6  
 incommensurability, 107–11, 113–15,  
 117, 119–21, 133, 125, 136–8,  
 194, 212, 253, 263, 270  
 infinitism, epistemological, 66–8, 246,  
 249, 254  
 internalism, epistemic, *see* evidentialism  
 invariantism  
 dogmatic, 171  
 sceptical, 171  
 subject-sensitive, 169–70, 172–5,  
 180, 188, 193, 211, 265–7, 276  
 traditional, 169, 172, 174, 193, 195,  
 211, 269  
 Jenkins, Carrie S.I., 19, 235–6, 239,  
 274  
 Jordan, James, 243  
 Joyce, Richard, 111  
 Kahane, Guy, 254  
 Kallestrup, Jesper, 193, 257–8, 265,  
 269–70  
 Kaplan, David, 53, 165, 167–8, 173,  
 261, 263–4  
 Kappel, Klemens, 226, 228, 230, 272  
 Keller, John A., 19  
 Kelly, Thomas, 102  
 Kelp, Christoph, 228, 268–9, 272  
 Klein, Peter, 61, 65–7, 246–8, 254  
 Knobe, Joshua, and Shaun Nichols,  
 111–13, 255  
 knowledge  
 analysis of (propositional), 4–8, 11,  
 60–1  
 ascriptions of, 62, 162, 164  
 basic, 61, 117–22  
 and epistemic luck, 21, 177–9,  
 182–3, 188, 193, 233, 265–9, 272,  
 276  
 factivity of, 11–12, 18, 196, 205,  
 207–10, 237, 273, 277–8  
 first approach, 18, 24, 197–9, 202,  
 205, 215, 264, 272  
 knowledge-how, 84–6, 104, 130, 251  
 perceptual, 112–14, 116, 118

knowledge – *continued*

- practical explications of, 226–32
- testimonial, 101, 189–91, 258, 269
- value of, 202–4
- virtue theoretic approach to, 92, 189–90, 200–1, 233, 235, 269, 272, 276
- Kölbel, Max, 32, 35, 37–8, 41, 105, 168–9, 240–2, 245
- Kuhn, Thomas, 63, 108, 116, 135, 257, 263
- Kupreeva, Inna, 245
- Kusch, Martin, 37, 137–9, 142–54, 156–7, 159–61, 194, 240, 245, 249, 255, 260–2, 270, 275, 277
- Kvanvig, Jonathan L., 202, 268, 272
  
- Lackey, Jennifer, 184, 255, 268–9
- Lammenranta, Markus, 61, 73, 75–6, 78, 101, 245–6, 250
- Latour, Bruno and Steve Woolgar, 263
- Lawlor, Krista, 228–9
- Lee, Mi-Kyoung, 42
- Lehrer, Keith, 11–12, 237–8
- Lewis, David, 53, 113, 167, 255, 265–8, 273
- Lipton, Peter, 198
- Littlejohn, Clayton, 102–3
- luck, epistemic, *see* knowledge, and epistemic luck
- Lycan, William G., 267
- Lynch, Michael P., 251, 253
  
- McCain, Kevin, 234, 247
- Maccallum, Gerald C., 124
- MacFarlane, John, 53, 55, 58, 105, 138–9, 156–7, 161–75, 178–83, 186–96, 200, 203–4, 209–11, 213, 220, 221, 224–6, 230–3, 236, 238, 241, 243, 250, 261–7, 269–70, 273, 278
- McGlynn, Aidan, 234, 238, 242, 257, 270, 272
- McKenna, Robin, 227–8, 251, 260, 263–9
- Madison, Brent, 267–8
- Maffie, James, 234
- Markie, Peter J., 259
  
- Matthen, Mohan, 243
- Meiland, Jack, 41, 242–3
- Miller, Alexander, 17, 19, 238–40, 274
- mind-independence, 14, 17–21, 24–5, 28, 33, 44, 50–1, 215–17, 238, 251, 257, 274
- Moore, George Edward, 3, 10–14, 19–20, 61, 125–7, 134, 152, 208, 223, 236–7, 239, 245, 264, 273–4
- Moore's Proof, 13–14, 125–6, 134, 236–7
- Mounce, H. O., 276
  
- Nagel, Thomas, 244, 262, 264
- Neta, Ram, 237, 257, 259
- Nozick, Robert, 34
  
- O'Grady, Paul, 240
- Olson, Jonas, 23, 26, 29, 218, 220, 275
- Olsson, Erik J., 246, 272
- openmindedness, 201–2
  
- paradigm shift, *see* Kuhn, Thomas
- Parfit, Derek, 233
- Passmore, John, 41
- Peacocke, Christopher, 272
- peritrope, 36, 37, 39, 46, 47, 49–50, 242, 244
- phenomenology, 86–7, 251
- platitudes, epistemic
  - authority, 22–3, 26–7, 29, 82, 217–21, 223, 240, 275–7
  - content, 22, 26–7, 29, 217–18, 221, 223, 239–40, 275–7
- Plato, 31, 36–40, 42–5, 47–51, 57, 240–3
- Pollock, John, 268
- presupposition
- pragmatic, 3, 9–10, 15, 28, 120, 152, 209–10, 222–3, 235–6, 256, 273
- semantic, 8–10, 236
- Pritchard, Duncan, 55, 76, 112–13, 120, 176, 193, 223, 227, 233, 235–6, 245, 247, 251–2, 255–6, 258, 260, 265–72, 276–7
- Proof, Moore's, *see* Moore's Proof
- Protagoras, 31, 36–48, 50–1, 54, 57–8, 155, 157, 240–4, 260–1

- Protagorean relativism, *see* truth, relativism about
- Pryor, James, 107, 125–7, 132–6, 237, 257, 259–60
- Pyrrhonism, *see* scepticism, Pyrrhonian
- realism  
 generic, 2, 3, 15–18, 44, 238–9, 274  
 metaepistemological, 1, 3, 7, 16, 18–25, 27–9, 30, 152, 212–18, 220, 223, 233, 238, 246  
 metaethical, 16, 152, 275
- reasons, epistemic, 22–4, 26, 218–20, 222, 275
- relationism, epistemic, 52, 140, 154–7, 159, 168
- Riggs, Wayne, 202, 227, 271–2
- Roberts, Robert, and Wood, Jay, 271
- Rorty, Richard, 27–8, 32, 74, 76–7, 79, 86, 166, 220, 263, 278
- Rothschild, Daniel, 236
- Ryle, Gilbert, 130, 260
- Rysiew, Patrick, 228, 263
- Salmon, Wesley, 271
- Sankey, Howard, 59–72, 75, 105, 245–9, 251, 253–4, 270
- Saul, Jennifer, 93–4, 252
- Sayre-McCord, Geoff, 1, 8, 25, 235
- scepticism, epistemological  
 Cartesian, 3, 11–14, 19–20, 36, 57, 61, 103, 152, 170–2, 175, 223, 233, 235–9, 245, 257, 269, 273–4  
 genealogical arguments for, 111–12  
 about inference rules, 130, 132–6  
 Pyrrhonian, 59, 61–3, 64, 66, 69–71, 73, 75–80, 88–90, 100–3, 105, 107, 138, 163, 194–5, 208, 212, 221, 245–50, 254, 270
- Schaffer, Jonathan, 268
- Schmitt, Frederick, 257
- science, Western, 28, 115, 123, 129, 140, 245, 257, 259
- Seidel, Markus, 69, 248–9
- self-refutation  
 contradictory, 37–8, 242, 262  
 dialectical, 37–40, 45–9, 51, 54, 57, 240, 242
- Shafer-Landau, Russ, 152, 233
- Shah, Nishi, 275
- Sider, Theodore, 43
- Siegel, Harvey, 76, 81, 84, 86–7, 104, 107
- situationism, 91–3, 95–7, 252
- Slovic, Paul, 252
- Sober, Elliott, 238
- Sosa, Ernest, 189, 201, 233, 254, 269
- Spencer, Jack, 34, 240–1
- Srinivasan, Amia, 111, 247, 254
- Stalnaker, Robert, 3, 8, 10, 12, 15, 110–11, 236, 273
- Stanley, Jason, 84–8, 197–8, 202, 205, 251, 265, 269, 270–2
- Strawson, Peter F., 8–9, 256
- Streumer, Bart, 240, 275
- Strevens, Michael, 271
- Stroud, Barry, 11, 12, 237
- styles of reasoning, *see* Hacking, Ian
- subjectivism, radical, 37
- Sudduth, Michael, 268
- swamping problem, 203, 272
- Swoyer, Chris, 33
- testimony, *see* knowledge, testimonial
- Toribio, Josefa, 84–8, 251
- truth  
 and knowledge, *see* knowledge, and factivity  
 nature of, 28, 32, 35–6, 38–59, 101, 103, 138, 155–8, 160, 166–7, 215, 238, 241–5, 260–5, 274, 278
- Turri, John, 65–7, 91–2, 207, 246–8, 258, 267, 269, 273–4, 278
- Tversky, Amos, and Daniel Kahneman, 92–3
- understanding  
 objectual, 272  
 why, 198, 204–5, 271
- Unger, Peter, 237
- virtue epistemology, *see* knowledge, virtue theoretic approach to
- Virvidakis, Stelios, 237
- Vlastos, Gregory, 242

- Vogel, Jonathan, 237, 257  
 Vogt, Katja, 250
- warrant, perceptual  
     conservatism, 107, 126, 128,  
         133–6, 236, 260, 270, 274  
     dogmatism, 107, 126, 128, 133–6,  
         171, 237, 246, 260, 270, 274  
     transmission of, 101, 122, 125–7,  
         130–1, 133–5, 258–60
- Williams, Michael, 114, 238, 246,  
     255–6, *see also* contextualism,  
     Williams-style
- Williamson, Timothy, 18, 24, 197–9,  
     205, 215, 219, 264–5, 269, 271–2
- Wittgenstein, Ludwig, 54–6, 58,  
     61, 74, 108, 115, 223, 237,  
     241, 245–6, 250, 256,  
     276–8
- Wittgenstenian hinges, *see* hinge  
     propositions
- Wong, David B., 149, 159, 255
- Woodward, James, 271
- Wright, Crispin, 23–5, 31, 34, 37, 46,  
     49, 52–3, 55, 57–8, 107, 126–7,  
     133–9, 153–61, 168, 237, 245,  
     248, 260, 263, 267
- Zagzebski, Linda, 189, 201
- Zande, *see* Azande, the