




OXFORD



# Social Epistemology

*edited by*

ADRIAN HADDOCK,  
ALAN MILLAR, AND  
DUNCAN PRITCHARD

# **SOCIAL EPISTEMOLOGY**

**OXFORD**  
UNIVERSITY PRESS



Great Clarendon Street, Oxford ox2 6dp

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide in Oxford New York  
Auckland Cape Town Dar es Salaam Hong Kong Karachi  
Kuala Lumpur Madrid Melbourne Mexico City Nairobi  
New Delhi Shanghai Taipei Toronto

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

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

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

© the several contributors 2010

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

First published 2010

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

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

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

ISBN 978-0-19-957747-7

1 3 5 7 9 10 8 6 4 2

# Contents

*List of Contributors* vii

*Acknowledgements* viii

*Introduction* ix

*Analytical Table of Contents* xi

1. Why Social Epistemology Is Real Epistemology 1

*Alvin I. Goldman*

2. Testimony, Advocacy, Ignorance: Thinking Ecologically about Social Knowledge 29

*Lorraine Code*

3. Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time 51

*Miranda Fricker*

4. On Saying that Someone Knows: Themes from Craig 69

*Klemens Kappel*

5. The Swamping Problem Redux: Pith and Gist 89

*Jonathan Kvanvig*

6. From Epistemic Expressivism to Epistemic Inferentialism 112

*Matthew Chrisman*

7. Norms of Trust 129

*Paul Faulkner*

8. Testimonial Entitlement and the Function of Comprehension 148

*Peter J. Graham*

9. Knowing from Being Told 175

*Alan Millar*

10. Can A Priori Entitlement be Preserved by Testimony? 194

*Ram Neta*

11. The Assurance View of Testimony 216

*Frederick F. Schmitt*

12. The Epistemology of Silence 243

*Sanford C. Goldberg*

13. Epistemic Circularity and Epistemic Incommensurability 262

*Michael P. Lynch*

14. The Epistemology of Disagreement 278

*Ernest Sosa*

15. A Justificationist View of Disagreement's Epistemic Significance 298

*Jennifer Lackey*

*Bibliography* 326

*Index* 345

## **List of Contributors**

**Matthew Chrisman**, University of Edinburgh

**Lorraine Code**, York University, Canada

**Paul Faulkner**, University of Sheffield

**Miranda Fricker**, Birkbeck College, University of London

**Sanford C. Goldberg**, Northwestern University

**Alvin I. Goldman**, Rutgers University

**Peter J. Graham**, University of California, Riverside

**Klemens Kappel**, University of Copenhagen

**Jonathan Kvanvig**, Baylor University

**Jennifer Lackey**, Northwestern University

**Michael P. Lynch**, University of Connecticut

**Alan Millar**, University of Stirling

**Ram Neta**, University of North Carolina, Chapel Hill

**Frederick F. Schmitt**, Indiana University

**Ernest Sosa**, Rutgers University

## **Acknowledgements**

We would like to extend warm thanks to the contributors to this volume for being willing to have their work published in it. Several contributions are related to presentations given at a conference on Social Epistemology held at the University of Stirling in late August-early September 2007. We are grateful to those who were present for the stimulus they provided. The conference was part of a project on the value of knowledge funded by the UK's Arts and Humanities Research Council. We thank the Council for its generous support.

Warm thanks are also due to Yvonne McClymont for assistance with the administration of the project, to Nikki Leavitt for secretarial assistance, to Steinvör Thöll Árnadóttir for assistance in preparing the typescript, and to Tom Cunningham for producing the Index. Their support has made the work of the editors very much easier.

Peter Momtchiloff has been an enormous support in relation to this and other published work emerging from the value of knowledge project. We very much appreciate his encouragement and assistance.

## Introduction

These are exciting times for the field of social epistemology, which has grown quite dramatically in the last twenty-five years to become one of the mainstays of contemporary philosophy. The aim of this volume is to offer a cross-section of the excellent work that is currently being conducted on this topic. We feel that we have succeeded.

The very idea of approaching epistemological concerns from a perspective that is both social and philosophical (as opposed to sociological) is relatively new. As has often been observed, the epistemological enterprise—indeed, arguably philosophy more generally—has for much of its history, and certainly since Descartes, been cast along egocentric lines. Where a non-egocentric approach has been taken, as in the recent work of naturalist epistemologists, the focus has been on individuals (not necessarily human) interacting with their environment rather than on the significance of social interaction for an understanding of the nature and value of knowledge. Indeed, until quite recently the only kind of social topic discussed at any length by mainstream epistemologists was that of testimony, and even then the issue tended to be largely understood from the perspective of the consumer of the testimony rather than in terms of the complex social web in which testimonial exchanges take place.

Things are very different today, with social epistemology now firmly lodged at the very heart of a vibrant contemporary epistemological literature. For example, consider the current discussion of testimony. From being a mainly peripheral part of the epistemological canon, testimony is now centre stage, and as part of this transition the very scope of the debate has altered to take in new topics.

This transition in the literature of testimony is well represented in the contributions to this volume. On the one hand, the papers by Ram Neta and Frederick F. Schmitt are specifically concerned with the work of a particular prominent figure in the literature (Tyler Burge and Richard Moran, respectively), and thus in this sense constitute natural extensions—and enrichments—of the mainstream discussion. On the other hand, the chapters by Jennifer Lackey, Michael P. Lynch, and Ernest Sosa explore a topic that is relatively new to the epistemology of testimony—the epistemology of disagreement—while other chapters take the mainstream debate about testimony in new directions by exploring the relevance of such notions as comprehension (Peter J. Graham), the social norm of trust (Paul Faulkner), and recognitional abilities (Alan Millar).



The real progress that has been made in the field of social epistemology is, however, best reflected in how the literature in this area now covers issues—often of a metaphilosophical flavour—which concern the topic of testimony only peripherally, if at all. This is reflected by the range of topics covered by the papers collected in this volume. At the metaphilosophical end of the spectrum there are the papers by Lorraine Code, Alvin Goldman, and Miranda Fricker which raise wide-ranging issues about the nature of social epistemology and what it can contribute to epistemology more generally. At the other end of the spectrum we have the chapters by Sanford C. Goldberg and Jonathan Kvanvig which examine particular epistemological issues that are salient to social epistemology (respectively, the epistemology of silence and epistemic value). Lying between these two camps are the chapters by Matthew Chrisman and Klemens Kappel, which are concerned with the specific meta-epistemological issue of the social dimension of knowledge ascriptions.

This volume thus captures the tremendous diversity and depth of the contemporary debate in social epistemology. In doing so our hope is that it advances that debate still further.

# Analytical Table of Contents

## 1. WHY SOCIAL EPISTEMOLOGY IS REAL EPISTEMOLOGY

How is social epistemology related to mainstream epistemology? Does it retain the traditional character of epistemology while merely giving it a social twist? Or does it advocate a radical socializing enterprise that heads down a very different path? In the latter case it might not merit the label ‘epistemology’ at all. In the former case, what exactly are the social twists and how do they relate to the mainstream? Three conceptions of social epistemology are distinguished here: *revisionist*, *preservationist*, and *expansionist*. Revisionism rejects mainstream assumptions, including the objectivity of truth and rationality, and it is plausible to deny it the status of ‘real’ epistemology. Preservationism is in keeping with mainstream epistemology and qualifies as ‘real’ epistemology. It studies epistemic decision-making by individual doxastic agents. What makes it social is its study of doxastic decision-making in the light of *social evidence*. Other preservationist topics include epistemic norms associated with various speech and communicational activities (assertion, debate, argumentation). Expansionism seeks to enlarge the reach of social epistemology while remaining continuous with the tradition. Its chief topics are the epistemic properties of collective doxastic agents and the influence of alternative social systems on epistemic outcomes. Illustrations of the social-system approach include (i) examination of legal adjudication systems and (ii) epistemic approaches to democracy. In law we can ask which of various trial systems (species of social epistemic systems) tend to generate the most accurate verdicts. In political theory democratic decision-making processes might be defended by appeal to their putative epistemic characteristics, for example, their reliability.

## 2. TESTIMONY, ADVOCACY, IGNORANCE: THINKING ECOLOGICALLY ABOUT SOCIAL KNOWLEDGE

Epistemologists in the twenty-first century come from a long tradition in which perception, memory, and testimony were viewed as the principal sources of knowledge. Of these, perception and memory, however enhanced, abstracted,

or elaborated, counted as the most reliable sources, with testimony ranking as a distant, and usually compromised, third. In my view, social epistemology reverses this ranking, granting a central place to testimony in the production of knowledge, and interrogating assumptions about the replicability and homogeneity of perception and memory. It thus generates a range of issues that had been *hors de question* for traditional epistemologists. Drawing on the conceptual framework I develop in my 2006 book, *Ecological Thinking: The Politics of Epistemic Location*, I propose that epistemic enquiry socially reconfigured is more fully naturalized than post-Quinean naturalized epistemology. Social epistemology focuses on epistemic practices communally engaged by identifiable knowers in the world (thus not principally in the laboratory), who are situated not just socially, but ethically, politically, demographically, geographically, and ecologically, where aspects of such 'situatedness' often count among the conditions that make knowledge possible. This enquiry focuses on testimony and advocacy as practices where these factors are particularly salient, and on ignorance not as a mere lack or failure of knowledge, but as a modality of not-knowing, or knowing inadequately/unjustly, which is itself situationally fostered, inhibited, or eradicated.

### 3. SCEPTICISM AND THE GENEALOGY OF KNOWLEDGE: SITUATING EPISTEMOLOGY IN TIME

The overarching purpose is to illustrate the philosophical fruitfulness of expanding epistemology, not only laterally across the social space of other epistemic subjects, but at the same time vertically in the temporal dimension. I set about this by first presenting central strands of Michael Williams' diagnostic engagement with scepticism, in which he crucially employs a Default and Challenge model of justification. I then develop three key aspects of Edward Craig's 'practical explication' of the concept of knowledge so that they may be seen to resonate positively with Williams' epistemological picture: the admixture of internalist and externalist features; the proto-contextualism; and, finally, the distinctively genealogical anti-sceptical impetus. In this way I aim to support and augment the socialized anti-sceptical case mounted by Williams, and so to show that expanding epistemology in the temporal dimension can be a productive move in central debates in epistemology.

#### 4. ON SAYING THAT SOMEONE KNOWS: THEMES FROM CRAIG

In his book *Knowledge and the State of Nature*, Edward Craig proposes a new methodological approach to one of the crucial questions in epistemology, the question what knowledge is, or more precisely the question regarding the nature and the concept of knowledge. The main aim of the chapter is to discuss and to propose a distilled version of Craig's approach. The main issue is how to make sense of questions such as 'What is the point of the concept of knowledge?' and 'Why do we attribute knowledge?' In the chapter, I characterize and defend a general way of understanding such questions, and contrast it to elements of Craig's approach. On the basis of this, I offer some views about what, in Craig's terminology, is the point of the concept of knowledge, and again I contrast this to some of the claims put forward by Craig. Like Craig, I believe that understanding the point of the concept of knowledge may yield some insights about knowledge and the concept of knowledge, though space does not permit a discussion of these further questions.

#### 5. THE SWAMPING PROBLEM REDUX: PITH AND GIST

The Swamping Problem is one of the central problems in the new value-driven approach to epistemology that has arisen recently. It arises from the fact that value is not always additive, so if you begin with something valuable (true belief) and add a further valuable thing (justification), there is no guarantee that the combination is more valuable. The usual target of such concerns is reliabilism, but such concerns plague approaches that are more conscious of value concerns as well, such as functional accounts that aim at an account of knowledge in terms of its social significance as a marker of dependable sources of information. Here I investigate the fundamental nature of the problem and several recent attempts to provide an escape route from the problem, concluding that none of them succeeds.

## 6. FROM EPISTEMIC EXPRESSIVISM TO EPISTEMIC INFERENCEALISM

Recent philosophical debate about the semantics of knowledge claims has largely centred on the question of whether epistemic claims are plausibly thought to be context sensitive. The default assumption has been that sentences that attribute knowledge or justification (or whatever else is epistemic) have stable truth conditions across different contexts of utterance, once any non-epistemic context sensitivity has been fixed. The contrary view is the contextualist view that such sentences do not have stable truth conditions but can vary according to the context of utterance. This debate manifestly presupposes that the meta-epistemological issue of accounting for the meaning of epistemic claims is to be settled by determining the truth conditions of these claims. In the light of two observations, I think this presupposition is under-motivated. First, many epistemologists see epistemic claims as evaluative or normative, in some sense. Second, in the meta-ethical debate most philosophers take alternatives to truth-conditional semantics, such as expressivism, as live options when it comes to evaluative or normative claims. As it turns out, I do not think expressivism provides a plausible account of normative concepts across the board. But considering it as an alternative in the meta-epistemological debate points the way to another alternative to truth-conditional semantics. This is a form of inferentialism. In this chapter, I try to motivate a move to epistemic inferentialism by showing how it overcomes worries about expressivism and interacts with plausible ideas about the social function epistemic claims play in our commerce with one another and the world.

## 7. NORMS OF TRUST

Should we tell other people the truth? Should we believe what other people tell us? This chapter argues that something like these norms of truth-telling and belief govern our production and receipt of testimony in conversational contexts. It then attempts to articulate these norms and determine their justification. More fully specified, these norms prescribe that speakers tell the truth informatively, or be trustworthy, and that audiences presume that speakers do this, or trust. These norms of trust, as norms of conversational cooperation, would then seem to be justified on the basis of the interest that each has in the cooperative outcome. The norms of trust would then be justified as Lewisian conventions. However, the joint outcome prescribed by these norms is not an equilibrium

point: a speaker always does better to have an audience's trust and the liberty to tell the truth or not, as it suits. In this way, testimony presents a problem of trust. The justification of these norms of trust then starts from the recognition that any society that did not resolve this problem of trust would be stymied as a society. The resolution of this problem then requires securing the motivations characteristic of trusting and being trustworthy, where to have these motivations is to have an ethical outlook defined in terms of internalizing these norms of trust. This justification is genealogical and it is one of value.

## 8. TESTIMONIAL ENTITLEMENT AND THE FUNCTION OF COMPREHENSION

Warrants epistemically support beliefs. There are two prominent kinds: entitlements and justifications. Justifications involve reasons and reasoning; paradigmatically justifications are arguments in support of a belief. Entitlements do not turn on reasons. Paradigmatically, warrant for perceptual belief turns on entitlements. It is argued that so-called 'testimony-based beliefs'—beliefs based on our capacity to comprehend assertive speech acts—enjoy entitlements, and do not rely exclusively for their warrant on justifications, when comprehension is accompanied by our capacity to filter out various kinds of unreliable or misleading assertions. For a belief enjoys *prima facie pro tanto* entitlement when based on a normally functioning belief-forming process that has forming and sustaining true beliefs reliably as a function. And comprehension-with-filtering has that function. Beliefs based on comprehension-with-filtering thus enjoy entitlement when our capacity to comprehend and filter functions normally, for our capacity to comprehend-and-filter has forming and sustaining true beliefs reliably as a function. The following topics are discussed: the psychology of our capacity to comprehend-and-filter; the nature of epistemic entitlement; the function of assertion; the function of our capacity to comprehend; the evolution of language; and the reliability of filtering strategies.

## 9. KNOWING FROM BEING TOLD

This discussion focuses on the transmission of knowledge through testimony in straightforward cases. These are cases in which subjects gain knowledge from the say so of a speaker without in any obvious way engaging in reasoning over the trustworthiness of the speaker. The account proposed accords a central

role to the practice of informing through telling, but while practice-theoretic considerations explain how it is that people who rely on testimony will so often acquire true beliefs, they do not explain how testimony can yield knowledge. We need to show that testimony that something is so can settle that it is so. The suggested account treats knowledge from testimony as a special case of knowledge through recognition of the significance of an indicator, for instance, a reading on a fuel gauge, or the presence of tracks on a path. To account for such knowledge we need to bring into play the notion of an ability to recognize the significance of an indicator—an ability that is honed by experience. These considerations are applied to the case of recognizing a person to be trustworthy in respect of an act of telling.

## 10. CAN A PRIORI ENTITLEMENT BE PRESERVED BY TESTIMONY?

In his influential article ‘Content Preservation’, Tyler Burge argues for the highly unorthodox thesis that, at least under some circumstances, when someone who is a priori entitled to believe that  $p$  tells us that  $p$ , she thereby renders us a priori entitled to believe that  $p$ . Anna-Sara Malmgren has argued against this thesis, and raised objection to what she takes to be Burge's argument for it. After devoting the first half of the present chapter to laying out precisely what Burge's thesis, and his argument, amount to, I show that Malmgren's argument against Burge rests on a misinterpretation. But I then go on to argue that, even when correctly interpreted, Burge's argument still does not work: it rests on the premises that (a) what you're entitled to believe is closed under modus ponens consequence, and (b) having an a priori entitlement to believe that you're entitled to believe that  $p$  suffices for having an a priori entitlement to believe that  $p$ . (a) and (b) appear plausible only if we confuse the externalist kind of warrant that Burge calls ‘entitlement’ with the internalist warrant that he calls ‘justification’.

## 11. THE ASSURANCE VIEW OF TESTIMONY

According to the *assurance view of testimony*, we are often testimonially justified in believing a proposition  $p$  in virtue of the testifier's assurance that  $p$  or the recipient's acceptance of that assurance. I focus here on the logically weakest assurance view, suggested by the work of Richard Moran, that such assurance or acceptance gives the recipient an epistemic reason to believe  $p$ , though whether

this reason is a good one depends on further 'background conditions'. On this view, testimonial reason-giving turns on an actual or potential personal relation between the testifier and the recipient and on their free choices to offer or accept an agreement. I object to this view, on three grounds. First, though assurance may give the recipient a practical reason to trust the testifier as to whether  $p$ , neither assurance nor accepting assurance is sufficient to give the recipient an *epistemic* reason to believe  $p$ , as opposed to a practical reason in the light of the cognitive goal of believing the truth. Second, assurance is not necessary for testimony to give the recipient an epistemic reason to believe the proposition. Third, the assurance view cannot account for the conferring of epistemic reasons by testimony in the most important class of cases for a non-reductive view of testimonial justification—the primitive cases in which the recipient lacks non-testimonial reasons. At the end of the chapter, I consider what these objections suggest for an alternative non-reductive view of testimony, the agreement view.

## 12. THE EPISTEMOLOGY OF SILENCE

In this chapter I explore an analogy between our reliance on memory and our reliance on news coverage. The analogy derives from a certain sort of counterfactual reasoning that might be employed to support beliefs in either domain, in cases involving the silence of a relied-upon source. Thus one might form the belief that  $p$  on the grounds that if  $p$  were false one would have heard about it by now (or would have remembered it). In another paper I develop an account of the conditions under which counterfactual reasoning of this sort is epistemically fruitful. Here I address a related matter: whether resources within the epistemology of memory might be helpful in thinking about the epistemic significance of silence from *any* relied-upon source—whether that source is one's own memory, or the news groups on which one relies for coverage.

## 13. EPISTEMIC CIRCULARITY AND EPISTEMIC INCOMMENSURABILITY

Epistemic disagreement is disagreement over epistemic principles, or principles concerning the reliability and extent of our epistemic methods. In this chapter, I argue that disagreement over this sort raises a new problem, distinct from scepticism. Like some sceptical arguments, the problem of epistemic incommensurability is rooted in part in the issue of epistemic circularity. But it is not a



problem about whether we in fact have knowledge or are justified in our opinions. It is about rationally resolving explicit disagreement over the reliability of our most basic methods for forming beliefs. The bulk of this chapter is concerned with getting clear on the problem and its nature. In the final section, I briefly sketch a solution.

#### 14. THE EPISTEMOLOGY OF DISAGREEMENT

This chapter takes up one main question: when and how can a belief be sustained reasonably in the face of known disagreement? On the answer to this question will depend our prospects for sustaining cherished opinions in fields where controversy abounds, such as religion, politics, history, morality, art, philosophy, and even medicine and the law.

#### 15. A JUSTIFICATIONIST VIEW OF DISAGREEMENT'S EPISTEMIC SIGNIFICANCE

In this chapter, I develop a justificationist account of the significance of disagreement between epistemic peers. Whereas current views maintain that disagreement, by itself, either simply does or does not possess epistemic power, my account holds that its epistemic power, or lack thereof, is explainable in terms of the degree of justified confidence with which the belief in question is held. In this sense, I reject nonconformism—the absence of doxastic revision in the face of peer disagreement is never justified merely by virtue of the fact that the beliefs in question are either mine or are the product of correct reasoning—and conformism—substantial doxastic revision in the face of peer disagreement is never justified merely by virtue of equal weight being given to my own beliefs and to those held by my epistemic peers. Despite this, however, one advantage of my justificationist account is that it is able to explain why nonconformism provides the intuitively correct result in some cases, while conformism gives the intuitively correct result in other cases. A further advantage is that my justificationist account is generalizable in a way that neither of these rival views is.

# 1. INTRODUCTION

Alvin I. Goldman

What is social epistemology? Or what might it be? According to one perspective, social epistemology is a branch of traditional epistemology that studies epistemic properties of individuals that arise from their relations to others, as well as epistemic properties of groups or social systems. A simple example (of the first sort) is the transmission of knowledge or justification from one person to another. Studying such interpersonal epistemic relations is a legitimate part of epistemology. A very different perspective would associate ‘social epistemology’ with movements in postmodernism, social studies of science, or cultural studies that aim to replace traditional epistemology with radically different questions, premises, or procedures. Although these enquiries examine the social contexts of belief and thought, they generally seek to debunk or reconfigure conventional epistemic concepts rather than illuminate the nature and conditions of epistemic success or failure. Under the first construal social epistemology is a bona fide part of the mainstream, and hence ‘real’ epistemology. Under the second construal, it is not part of epistemology at all. If protagonists of the latter-day movements marketed their products under the guise of epistemology, they would be imposters. Their products are not real epistemology.

‘Is social epistemology *real* epistemology?’ is a question posed by William Alston (2005). He raises it *en passant*, while noting that epistemology’s boundaries are controversial, drawn differently by different thinkers. To illustrate his point, he suggests that much of the material in my book on social epistemology, *Knowledge in a Social World* (Goldman 1999), ‘would be rejected by many contemporary epistemologists as “not real epistemology”’ (Alston 2005: 5). The epistemologists in question, says Alston, would relegate much of this so-called social epistemology to sociology, social psychology, or other social sciences, or perhaps to the philosophical foundations thereof.

1

What about this claim that Alston attributes to imagined readers of *Knowledge in a Social World*? Would these readers be correct in saying that the material in question does not qualify as ‘real’ epistemology? Setting aside the specifics of *Knowledge in a Social World*, which theoretical enterprises that fly—or might fly—under the flag of ‘social epistemology’ deserve the label ‘epistemology’ and which ones do not? Wherein does social epistemology differ or diverge from traditional epistemology, and what differences and divergences would carry social epistemology beyond the pale?

## 2. THREE CONCEPTIONS OF SOCIAL EPISTEMOLOGY

I shall divide conceptions of social epistemology (SE) into three types: (1) *revisionism*, (2) *preservationism*, and (3) *expansionism*. These conceptions will be framed in terms of their relationships to traditional or mainstream epistemology. I shall agree that revisionist SE should not be classified as ‘real’ epistemology. By contrast, both preservationist and expansionist forms of SE deserve to be considered bona fide branches of epistemology. The difference between the two is that preservationist SE clings more closely to traditional paradigms of mainstream epistemology whereas expansionist SE ventures a bit further into unfamiliar terrain.<sup>1</sup>

As a prolegomenon to an assessment of SE's prospects, then, let us consider central features of epistemology ‘as we know it’ (in the analytic tradition). Here is a catalogue of such features. (A) The epistemic agents of traditional epistemology are exclusively individuals. (B) Epistemology focuses on the study of epistemic evaluation or normativity, represented by such evaluative concepts as justifiedness, rationality, and knowledge. Traditional epistemology asks such questions as how individuals can acquire knowledge and maintain justified or rational credal states. (C) Traditional epistemology assumes that the normative standards of rationality and justifiedness are not merely conventional or relativistic, but have some sort of objective validity. (D) The central notions of epistemic attainment—knowledge and justification, for example—either entail truth or are closely linked to truth. A known proposition must be true; and justified beliefs, according to many mainstream views, are in some sense likely to be true. (E) Truth is assumed to be an objective, largely mind-independent, affair. (F) Traditional epistemology takes its central business to be the critical examination of doxastic ‘decision-making’ (DDM): adopting, retaining, or revising one's beliefs and other doxastic attitudes. Let us call the foregoing assumptions the *core* assumptions

2

of mainstream epistemology. The three conceptions of social epistemology can then be elucidated and distinguished by their different postures vis-à-vis these core assumptions.

### 3. REVISIONISM

The writers subsumed under the heading of revisionism rarely employ the label ‘social epistemology.’<sup>2</sup> Nonetheless, since they typically study science and other forms of ‘knowledge’-seeking, and they certainly emphasize the social character of these pursuits, mainstream philosophers might view their approaches as aspiring forms of epistemology. Under revisionism I include postmodernism, deconstructionism, social constructionism, and various social studies of science, including the ‘strong programme’ in sociology of science. These movements share a staunch rejection of many if not all of the core tenets

of traditional epistemology. These rejections, moreover, are inspired by the social character of the phenomena they study.

A few philosophers such as Richard Rorty fall into the same category, sometimes waxing emphatic about the bankruptcy of traditional epistemology. Rorty (1979) declared the 'death' of epistemology and proposed a vague replacement for it in the form of a 'conversation of mankind'. His mantra was to 'keep the conversation going rather than to find objective truth' (1979: 377).

One strand of social constructionism holds that truths or facts are not in or of the world; they are not 'out there' to be discovered but are mere social fabrications or constructions (Latour and Woolgar 1986). As Steven Shapin puts it, 'truth is a social institution' (1994: 6). While revisionists often use the language of 'knowledge,' they don't understand it to be a truth-entailing, or factive, state. In their lexicon knowledge is simply whatever is believed, or perhaps 'institutionalized' belief. On the topic of rationality, we find proponents of the strong programme rejecting objectivist rationality, declaring that 'there are no context-free or super-cultural norms of rationality' (Barnes and Bloor 1982: 27). On the question of epistemic agents, at least a few revisionists deny that individuals are epistemic agents at all. Only groups or communities qualify as knowers, according to Lynn Hankinson-Nelson (1993):

In suggesting that it is communities that construct and acquire knowledge, I do not mean . . . that what comes to be recognized or 'certified' as knowledge is the result of collaboration between . . . individuals who, *as individuals*, know in some logically or empirically 'prior' sense . . . [T]he collaborators, the consensus achievers, and, in more general terms, the agents who generate knowledge are communities and subcommunities, not individuals. (Hankinson-Nelson 1993: 124)

3

Thus, revisionists dispute many of the main tenets of traditional epistemology. They might be called 'social epistemologists' by courtesy, but there is a wide gulf between their questions and those that traditionally go by the name 'epistemology'. Researchers in these traditions are often absorbed with case studies in history and/or sociology of science, with only scant attention to epistemic implications as mainstream epistemologists would conceive of them. Several writers in the sociology of science advance political-military accounts of science, aiming to displace the traditional notion that science is a province of reason or rationality. Shapin and Schaffer remark that in scientific disputes, '[h]e who has the most, and the most powerful, allies wins' (1985: 342). Latour (1987: 31) offers a detailed analysis of how scientific controversies in science are conducted by enlisting the support of forceful 'allies' such as prominent journals, Nobel prize winners, or granting agencies. What does this show about the *epistemic* properties of science? Latour's political-military account is intended to debunk or undercut the alleged role of evidence and truth in science.<sup>3</sup> This makes it natural for

mainstream epistemologists to deny that such approaches constitute genuine epistemology, because little or no serious epistemic business gets conducted, as mainstreamers see it.<sup>4</sup>

Our attempt to draw a firm boundary between revisionism and mainstream epistemology may be a bit quick. Recent developments in epistemology and philosophy of language may argue for greater nuance. Is it true, for example, that mainstream epistemology firmly and definitively rejects relativism? Recently the scene has shifted a bit. There is a ‘New Age’ kind of relativism—as Crispin Wright (2007, 2009) calls it—that has gained popularity within analytic philosophy. Perhaps New Age relativism has much in common with what revisionists like Barnes and Bloor have claimed. If so, there may be less distance between revisionism and mainstream epistemology than suggested above.

However, we should not exaggerate the degree to which the popularity of New Age relativism closes the gap between revisionism and mainstream analytic epistemology. What Wright calls ‘New Age’ relativism is a view championed by John MacFarlane (2005), Max Kolbel (2003), and Peter Lasersohn (2005), among others, that says the following: for a number of regions of discourse

4

the very same token claim can take different truth values when considered in different *contexts of assessment*. This doctrine is a version of contextualism, except in holding that an utterance's truth value depends on an index that varies with variable characteristics of hypothetical *assessments* (rather than, say, contexts of *use*). In discussing predicates of personal taste, Lasersohn suggests (at first pass) that in analysing a sentence like ‘Roller coasters are fun’, we let the *character* (in Kaplan's sense) of the predicate ‘fun’ yield different *contents* in different contexts, depending on the speaker (or some other referent). It might be claimed that the predicate ‘fun’ has a hidden argument place, filled by some sort of implicit indexical pronoun. Carrying this over to epistemology, New Age relativism would be the thesis that the truth of statements about epistemic justification is assessment-relative, with systems of standards providing the relevant assessment-contextual parameter (Wright 2009: 384). This kind of relativism is attracting non-negligible support among analytic philosophers.

However, as Paul Boghossian points out (2008: 414), New Age relativism is primarily a *semantical* thesis, and as formulated above it does not express everything—or the most important thing—that revisionist relativism means to assert. Revisionist relativism would agree on a central presupposition of New Age relativism, namely, that multiple epistemic systems are possible and justification claims can be assessed relative to any of these different systems. Revisionist relativism, however, also endorses a further thesis, which Boghossian sometimes calls ‘epistemic nihilism’. Epistemic nihilism says that ‘there are no facts by virtue of which one of these systems is more correct than any of the others’ (Boghossian 2006: 73, 2008: 413).<sup>5</sup> This, I submit, is the core thesis of relativism that

revisionists advance and mainstream epistemologists dispute. So, even conceding the respectability of New Age relativism, there remains a significant divide between revisionist and mainstream views in the matter of relativism (in addition to other core differences).

In short, with respect to a body of work in the social sciences and the humanities that is sometimes dubbed ‘social epistemology’, epistemologists are well within their rights to deny it the status of *real* epistemology. If one is in the demarcation business, there are reasonable boundaries to be drawn.

#### 4. PRESERVATIONISM

Nonetheless, many questions that currently preoccupy mainstream epistemologists or that interested historical epistemologists have an indisputable social

5

dimension.<sup>6</sup> Suitable treatments of these questions would preserve the core assumptions of traditional epistemology. So if SE is equated with the sector of epistemology that scrutinizes *these* questions, SE is well within the realm of legitimate, genuine epistemology. I shall illustrate preservationism by reference to three types of activity studied by traditional epistemology, each of which has a social sector. Each social sector of these parts of epistemology qualifies as ‘real’ epistemology. The three types of activity, in their social variants, are (1) doxastic decision-making with social evidence, (2) gathering social evidence, and (3) speech and communication with an informational purport (assertion, debate, argumentation, etc.).

##### 4.1. Doxastic Decision-Making with Social Evidence

As we have said, the central business of traditional epistemology is the epistemic evaluation of doxastic decision-making (DDM). In general there is nothing social about DDM. Belief-forming and belief-revising processes are not themselves inherently social, and evidence used as a basis for forming and revising beliefs need not involve subject matter (i.e. content) concerning other people either. But one sector of DDM is squarely based on the use of social evidence. Here the contents of the evidential beliefs do concern other people. First, the evidential beliefs may have contents concerning what other people have said or written. Second, the evidential beliefs may have contents concerning other people's opinions (or other psychological states), which the doxastic agent may acquire at second or third hand. In either type of case the doxastic agent uses social evidence.<sup>7</sup>

Two lively portions of current epistemology, *testimony* and *peer disagreement*, focus on the use of social evidence in these ways. Testimony poses the problem: ‘When is a doxastic agent justified in accepting another person’s testimony?’ This is arguably the fundamental topic of social epistemology. It occupied such historical philosophers as Hume (2000) and Reid (1983), and occupies a prominent place in the contemporary literature on epistemic justification

6

(Coady 1992; Fricker 1995; Lackey and Sosa 2006). Peer disagreement raises the question of whether rationality requires one to revise one’s belief (or degree of belief) if one finds oneself in disagreement with someone else who shares roughly the same evidence and has comparable cognitive abilities.<sup>8</sup> Each topic concerns epistemic justification or reasonability, and thereby falls within mainstream epistemology. Moreover, each features social evidence and therefore qualifies as a problem of social epistemology (SE). Hence, *some* social epistemological problems belong to the mainstream agenda. Some parts of SE unambiguously qualify as ‘real’ epistemology—real because they preserve or exemplify the tradition.

Notice that although different epistemologists take different stances in the dispute over peer disagreement, they all appeal to DDM principles to defend their positions. Those who advocate ‘splitting the difference’ with one’s peer (Christensen 2007; Elga 2007) tend to appeal to Bayesian principles, whereas non-splitters tend to appeal to other epistemic principles (Kelly forthcoming; Sosa this volume; Lackey this volume; Feldman and Warfield 2010). There is no mystery, then, about how social and traditional epistemology can comfortably intersect. Roughly similar SE work in a preservationist mold antecedes the current debate over peer disagreement. This includes Lehrer and Wagner’s (1981) model of rational consensus, in which agents assign levels of trust to self and others and update their opinions by taking account of others’ opinions, weighted by the degrees of trust they are accorded. This model also treats agents as consumers of social evidence.

#### **4.2. Gathering Social Evidence**

Another important epistemic activity in addition to DDM is deciding whether and how to gather *new* evidence. Such evidence gathering is variously called ‘enquiry’ or ‘investigation.’ In science it includes the design and implementation of tests, measurements, and experiments. Recent epistemology, narrowly construed, does not display extensive discussion of the theory of evidence gathering, investigation, or experiment. But there is plenty of discussion in the philosophy-of-science branch of epistemology. In a memorable phase of philosophy of science Karl Popper (1962) defended a falsificationist approach to scientific experimentation over a verificationist approach. Another sample admonition in the theory of testing, endorsed by assorted philosophers of science, is the desideratum of seeking a *variety* of evidence. This maxim flows from the thesis, shared by many confirmation theorists, that more varied evidence lends greater confirmation to a hypothesis than less varied evidence. A more recent

tradition in the theory of evidence is embedded in the Bayesian approach, especially Bayesian decision

7

theory. I. J. Good (1967) explained the value of gathering new evidence in terms of increased (subjectively) expected utility, assuming that the cost of getting such evidence is negligible (cf. Skyrms 1990).

Here is another toy principle based on Bayesian principles, but where the assumptions about probability are not exclusively subjectivist and the expected benefits are purely epistemic. If it is known that one possible test that could provide evidence on a given question has a more extreme likelihood ratio than an alternative test, the more probative and hence choice-worthy test is the former, more discriminating, test. At least it is preferable if the agent uses conditionalization to update his probability assignments to the target hypothesis. This principle can be rationalized in terms of a 'truth acquisitional' goal if certain assumptions are made about the measurement of truth possession and the accuracy of the agent's subjective likelihoods. A greater (objectively) expected increase in truth possession is associated with a more discriminating test than with a less discriminating test.<sup>9</sup>

To illustrate the last-mentioned result, suppose you want to gather more evidence about whether it will rain tomorrow. Two possible tests are available to you: listening to a weather report and flipping a coin. Concerning the weather-report test, suppose you antecedently estimate the likelihood of hearing a rain prediction on the weather report given that it is going to rain as .90, and you estimate the likelihood of hearing a rain prediction given that it is *not* going to rain as .15. Then the likelihood ratio of hearing a rain prediction using the weather report test is  $6:1 = 6$ . Turning to the coin-flipping test, assume that your subjective estimate of the likelihood of heads given that it is going to rain is .50, and you make the same estimate of the likelihood of heads given that it isn't going to rain. Then the likelihood ratio of the heads outcome for the coin toss is 1. Obviously, the likelihood ratio for the weather report test is more extreme than for the coin-flipping test. Hence, according to the result reported above, if all of your likelihood estimates are *accurate* and you apply (Bayesian) conditionalization, there is a larger expected increase in truth-possession from the weather-report test than from the coin-tossing test. (The requirement of accuracy, of course, presupposes both subjective and objective likelihoods.) This comparison of the two tests in terms of their expected increase in truth-possession holds no matter what your *prior* subjective degree of belief may be vis-à-vis rain (assuming it is neither 0 nor 1). The commonsensical upshot is that not all tests are equal in terms of their prospects for enhancing degree of truth possession.

8

Despite these examples, there are epistemologists who deny that evidence gathering is a proper subject for epistemological assessment. Richard Feldman (2000), in particular,



resists this idea. He seems to hold that the only fundamental epistemic goal is to have reasonable beliefs, and this does not require evidence gathering.

If the fundamental epistemic goal is just to have reasonable beliefs, then nothing about evidence gathering techniques or the like follows as a means to that goal. This goal just has implications concerning what attitudes we ought to take given the evidence we have. By seeking out new evidence concerning some important proposition and then believing what that evidence supports, I don't do a better job of achieving the goal of believing reasonably. I achieve that goal at any moment by believing what is then supported by my evidence. It's surely true that there are times when one would be best off finding new evidence. But this always turns on what opinions one has, what one cares about, and other non-epistemic factors. As I see it, these are prudential or moral matters, not strictly epistemic matters. (Feldman 2000: 188–9)

What Feldman means by 'reasonable belief' seems to be reasonable *degree of belief*, where this can range from a strongly confident belief to complete agnosticism. Either type of state is as good as the other, if both are reasonable. This implies that avoiding or getting rid of evidence does not hurt one's evidential state. Getting rid of evidence implies that one must choose a different degree of belief, but one can still be as reasonable as before, and this is all that counts, according to Feldman.

The evidentialist principles stated here imply that one can get oneself into a highly rational state by ridding oneself of as much evidence as one can and then suspending judgment about virtually everything that comes to mind . . . . By believing very little, he'll then be highly rational according to evidentialist standards . . . . But that seems to be exactly the right conclusion. Once the person has lost his evidence, he has no reasons to believe much, and he'd be unreasonable if he did believe things that would have been well-justified for him had he been in more normal circumstances. (2000: 189–90)

To me, however, this indicates that Feldman's account of our fundamental epistemic goals is inadequate. It is easiest to show this by considering a case where one's *practical* interest requires something more than just reasonable degree of belief (understood as being satisfiable by reasonable agnosticism). The practical goal is best attained by an epistemic state that consists of either *true* (categorical) *belief* or a high degree-of-belief (HDOB) in a truth. Since *true* belief or HDOB is unlikely to be attained without having good reason or justification for it, the appropriate epistemically valuable state to obtain is a reasonable *belief* or *HDOB*, not simply reasonable agnosticism.

To illustrate, suppose that on my return from a shopping trip, I discover that my credit card is missing, and I evidently left it at one of two stores where I used the card while shopping. Where did I leave it? At the moment, I am agnostic about where I left it; that is, I am agnostic with respect to each of the propositions

'The card is at store X' and 'The card is at store Y'. Suppose that the two stores are at opposite ends of town, so I will avoid wasting my time if I make a trip to only one of

them, not both. This can reliably be achieved if I form a true belief (or true HDOB) about where the card is. It cannot be reliably achieved if I remain agnostic about where the card is. (If I have to flip a coin to decide which store to return to first, there will be only a 50–50 chance of choosing the right one; so chances are even that I will have to return to both. On the other hand, if I get a true belief or true HDOB, I won't need to flip a coin; I can go to the right store straightaway and achieve the practical goal: retrieval of the card.) Of course, the way to obtain a true belief (or HDOB) of the requisite sort is to obtain a reasonable belief of this kind, assuming that reasonable beliefs are generally true. To form a reasonable belief about where I left the credit card, what I clearly need to do is gather some pertinent evidence. For example, I should telephone the two stores and inquire about my lost credit card. As this case shows, practical interests commonly dictate the desirability of having epistemic states that are not merely reasonable degrees-of-belief (DOBs) but reasonable *high* DOBs (HDOBs). And getting such epistemic states commonly requires gathering evidence. Not just gathering any old evidence, but gathering *enough* and *suitable* evidence. It is appropriate, therefore, for epistemology to counsel epistemic agents in the execution of this ubiquitous task. (Bad methods of evidence gathering, such as confirmation bias, should be avoided.)

What of Feldman's point that whether and how to get new evidence always turns on prudential or moral considerations, hence are not 'strictly' epistemic questions? He writes:

There are cases in which one can spend one's time gathering evidence about propositions concerning inconsequential and trivial propositions or about more weighty matters. Evidentialism provides no guidance about what to do. As I see it, this is not a weakness of evidentialism, since such choices are not to be made on epistemic grounds. (2000: 189)

Two comments are pertinent here. Undoubtedly, which topics deserve evidence-gathering attention or energy is often a purely practical matter, not an epistemic one. Nonetheless, *given* that topic T deserves attention or energy, questions about which kinds of evidence are relevant (and maximally helpful) and how to gather such evidence come into play. These choices are squarely epistemic, at least largely so. Second, even if a residue of the practical infuses these decisions (e.g. how much money and time to devote to evidence gathering), recent epistemology has abundantly shown that epistemic matters in general cannot be divorced from the practical. In the analysis of both knowledge and justification, proponents of the 'pragmatic encroachment' thesis have persuasively argued that standards of epistemic assessment are raised and lowered as a function of practical interests (DeRose 1995; Stanley 2005; Fantl and McGrath 2002). Finding that similar pragmatic factors play a role in evidence gathering should not remove this type of activity from the epistemic sphere.

We have not said anything yet about *social* evidence gathering specifically. This is a special province of SE under my proposal. It is not essential to SE's viability that there be wholly distinct principles of evidence gathering (or doxastic decision-making) for the

social domain. There may be only general principles of evidence gathering that apply equally to social and non-social domains. This would not diminish the interest of their application to the social-evidence domain.

However, there may well be distinctive aspects of evidence gathering in the social domain. One is the choice among experts (or putative experts) from whom to solicit opinions. Sometimes one would like to know what the world's leading experts think on a given subject. Sometimes one might be content with getting an opinion from a reasonably competent expert, for example, one's personal doctor or dentist. In both cases, however, the question arises of who meets the specified condition. This is often difficult to determine, especially if the field is a relatively arcane one and the evidence-seeker is a layperson who is largely ignorant of the field (Goldman 2001). Since such a person has no (confident) opinions of his own in the area, he cannot determine various candidates' expertise by comparing their answers to his own in order to measure their accuracy. Should he seek meta-experts to tell him who are the experts? How shall he identify the meta-experts? This is a sample problem specific to SE, which can be viewed as a problem of social evidence gathering.

Special problems in the area of social evidence gathering may also be connected with issues that arise in the peer disagreement literature. In judging whether or to what extent another person is a peer (that is, has the same intellectual competences as oneself), is it legitimate to rely on one's own first-order beliefs to determine the other's competence—and hence trustworthiness as a source to employ for gaining *additional* evidence? If one uses same-mindedness as a criterion of trustworthiness, doesn't one run the risk of consulting those who share the same biases, false theories, or bad company that plague one's own first-order belief set? Does this mean that one should set aside one's first-order beliefs in deciding who are reliable sources of further evidence and turn instead to higher-order evidence about belief-forming methods and sources? Some epistemologists argue that higher-order evidence brings special problems of its own into the picture (Kelly forthcoming; Christensen forthcoming).

### **4.3. The Social Epistemology of Speech and Communication**

Belief formation and evidence gathering are the most commonly discussed epistemic activities but they don't exhaust the field. Another type of social-epistemic activity is making statements and engaging in argumentative exchange. Such speech activities are social, of course, and the analysis of their epistemic properties is naturally assigned to SE.

A prominent topic in this corner of current epistemology is the question, 'What is the norm of assertion?' Specifically, what epistemic (or psychological) condition

must one satisfy to be entitled to assert something (to others)? Knowledge is one candidate answer,<sup>10</sup> and others include true belief, justified belief, adequate evidence, and so forth. What kind of norm is sought here? It is assumed that it is an epistemic norm, a

norm that 'comes' from epistemology rather than ethics or conversational practice, for example. But there is no incompatibility in saying that the source of the norm straddles two categories. A Gricean (1989) perspective suggests that some norms of speech arise from a cooperative enterprise for the social sharing of information (see also Craig 1990; Williams 2002). This cooperative enterprise is, effectively, a social-epistemic enterprise, so one can view the norms as having both a linguistic and an epistemological origin.

Turning from the simple act of assertion, consider more complex speech activities that invite epistemological attention, especially argumentative exchange. Argumentation and dialectic have long been part of the epistemological mix, though often more at the periphery than the centre. Sometimes this topic has been treated under the heading of logic rather than epistemology. Our interest, of course, is its suitability under the latter heading.

A successful dialectical approach to justification would position a social activity at a centre of epistemological analysis. One gets glimmers or whispers of such an approach in a variety of epistemological writings. But such an approach has not been thoroughly developed, and I myself think that the prospects are gloomy. Nonetheless, because it *would* be very important for SE if such an approach succeeded, it is worth assessing the prospects.

A purely dialectical approach to justification might say that a belief of S is justified if and only if S can satisfactorily respond to critics or challengers, to their requests or demands for reasons or evidence. Approaches to justification of this ilk appear in the work of David Annis (1978), Keith Lehrer (1990), and Michael Williams (2001).<sup>11</sup> It also has Wittgensteinian (1969) resonances. None of these writers, however, specifically labels his approach a 'dialectical' theory. Annis and Williams favour the term 'contextualism'. Alston (1989: 236), while not endorsing a dialectical analysis of justification, argues that our justification concept developed because of the practice of challenging one another's belief credentials and responding to such challenges.

How might a dialectical analysis of justifiedness go? To begin with, we should clarify the point that the analysandum, *being justified*, is a property of doxastic states or doxastic agents. Call this property *personal justification*. According to the envisaged approach, it is to be analysed in terms of a possessor's disposition to interact verbally with others in some specified fashion. The requisite verbal

12

interaction may be called *interpersonal justifying*, something one does *to* or *with* interlocutors. A sample dialectical analysis might run as follows:

(IPJ) S is personally justified in believing P if and only if S can or would succeed in *interpersonally justifying* P to (actual or potential) challengers.

IPJ isn't a bald persuasion theory. S's being justified in believing P, according to IPJ, does not require that S could persuade (erstwhile) challengers of P to believe it. A persuasion theory would be both too strong and too weak. Too strong because challengers might be

too stubborn or thick-headed to be persuaded even when offered robust evidence. Such intransigence should not keep S from being justified. Too weak because challengers might accede to a speaker's claim that P through excessive gullibility or deference. Such responses would not confer genuine justification on S's belief.

IPJ does not confront these liabilities, but it still faces hurdles. What does it mean to justify a proposition *to* another person? Is interpersonal justification to be defined in terms of personal justification, for example, as 'bringing it about (by verbal means) that the hearer is (personally) justified'? That would introduce a blatant circularity into the approach. One might try to offer a non-circular definition of interpersonal justification, but exactly how this would go is unclear. The definition would have to say things about deductive or inductive relationships between the asserted proposition (the 'conclusion') and the reasons the speaker might offer in its support. It would also have to say things about the status of those reasons or premises *for* the hearers or interlocutors. What is the requisite status? This is where things get murky. Moreover, there are problems of informational asymmetries between speaker and hearers. Suppose the speaker possesses (undefeated) counterevidence to his conclusion but doesn't disclose this evidence to the hearers. The hearers may be justified in believing P by what the speaker says, while the speaker himself is (personally) *un*justified, in virtue of the defeaters he possesses.

These reflections make me dubious of the viability of the dialectical approach to being justified. It isn't *this* easy to establish social factors at the centre of the epistemic realm. However, there are neighbouring topics in epistemology to which speech and argumentation are germane. The philosophical study of dialectic has a long and substantial history, often subsumed under the heading of logic, or an autonomous field called 'informal logic'. I propose instead that the normative theory of dialogue, dialectic, debate, or argumentation be viewed as part of social epistemology.<sup>12</sup> Most contributors to this literature do not offer

explicitly epistemological perspectives, but it is the most fruitful one, I submit (Goldman 1994, 1999: ch. 5, 2003).

Here is a simple illustration of how a social epistemological approach to argumentation helps make sense of the norms of interpersonal reasoning. A principal task in this field is to explain which forms of argumentation are in some sense legitimate and which are not. Argumentative legitimacy is not equivalent to validity (or inductive strength). Indeed, a prominent fact that needs explaining is that *enthymemes*, which are invalid arguments as they stand, are commonly used in everyday speech; yet nobody finds them objectionable. What sense can be made of this? It does not suffice to observe that enthymemes can be turned into valid arguments by supplying additional premises. That's true, of course, but not every enthymeme would be acceptable in real discourse. Acceptable ones must bear other suitable properties. Which properties? A straightforward account and explanation is an epistemic one (Goldman 2003: 60). Suppose that the usual aim of argumentative discourse is to justify a proposition to one's intended audience, that is, change the

epistemic state of the hearers so that they are (newly) justified in believing the conclusion. Assume further that in assessing the plausibility of a proffered conclusion, a hearer will not rely exclusively on explicitly stated premises. She will marshal all relevant evidence at her disposal, at least evidence that readily comes to mind. Finally, suppose that on a given occasion evidence E is relevant to the speaker's conclusion, and E would be perfectly appropriate to include as a premise in his argument. As the speaker realizes, however, E is already widely known in the discursive community. Then the speaker is entitled to assume that his hearers already know E and would think of it in contemplating his proffered conclusion. Why, then, bother to mention E? Conversational efficiency dictates otherwise, and epistemic considerations concur. A hearer will use E to help draw the same conclusion whether the speaker mentions it or not, and his *justifiedness* vis-à-vis the conclusion will be the same. Thus, presentation of the enthymematic argument, with E left unexpressed, is entirely in order. But this is so because of the audience's epistemic circumstances and their interrelationship with the speaker's epistemic end. Thus, an otherwise puzzling feature of discourse is nicely illuminated when viewed from a social-epistemic perspective.

In summary, many forms of SE are unquestionably instances of traditional epistemology. It would be arbitrary and misguided of epistemologists to contend that *no* part or instance of SE is real epistemology.<sup>13</sup> However, if SE were restricted to the two topics identified earlier—testimony and peer disagreement—or even if we add to SE's agenda the other tasks sketched in this section, many *other* projects described in *Knowledge in a Social World* would remain beyond the pale.

## 14

They would still not qualify as 'real' epistemology. Thus, restricting SE to the problems and projects that have clear precedents in historical and mainstream epistemology might still be unnecessarily limiting. In the rest of the chapter I argue for extending it beyond these confines.

### 5. EXPANSIONIST SOCIAL EPISTEMOLOGY

I argue for expanding SE's scope by adding (at least) two or three further topics to its agenda. These are not explicitly addressed in traditional epistemology but are sufficiently continuous with traditional epistemology that they deserve inclusion under the heading of 'real' epistemology. The topics are, first, the epistemic properties of group (or collective) doxastic agents and, second, the influence of social 'systems' and their policies on epistemic outcomes.

#### 5.1. Collective Doxastic Agents

Numerous philosophers in recent years have endorsed the idea of recognizing group or collective entities as subjects of intentional attitudes (Gilbert 1989; Bratman 1993; Tuomela 1990; Schmitt 1994a; Pettit 2003; List and Pettit forthcoming). Many of these philosophers are primarily interested either in the ontology of group agents or in aspects of practical reasoning and group commitment. Intentions have tended to be their favoured

intentional attitude. But there are equally good reasons to include attitudes of affirmation and denial of factual propositions, that is, judging propositions to be true or false, as properties of collective subjects. These are attitudes with a mind-to-world direction of fit rather than a world-to-mind direction of fit. In everyday life we often speak of collective entities like committees, juries, courts, scientific panels, and governments as making judgements, or taking cognitive stances, of one sort or another. Once these kinds of intentional attitudes for collective agents are conceded, shouldn't we reflect on the epistemic properties of these attitudes? If an individual's judgements over a set of many propositions can be assessed for their (epistemic) rationality, isn't it appropriate to make similar assessments of a collective agent's judgements? And if epistemologists take an interest in cases that seriously challenge what is rational for the individual, shouldn't epistemologists—social epistemologists—take an interest in cases that seriously challenge what is rational for a collective agent? As we shall see, challenges to rationality (and related normative desiderata) of a collective agent's judgements are, if anything, even more extensive and sobering than challenges to an individual agent's judgements.<sup>14</sup>

15

Worries about collective judgemental rationality were first identified by Kornhauser and Sager (1993) in the context of courts of law. They called this challenge 'the doctrinal paradox'. Suppose that a three-member court has to render a judgement in a particular breach-of-contract case, and legal doctrine specifies that a defendant is liable for breach of contract if and only if he was legally obliged not to do a certain action but nonetheless did it. Then the court must render judgments on each of several propositions concerning the particular case:

- \* The defendant was legally obliged not to do a certain action (first premise)
- \* The defendant did the action (second premise).
- \* The defendant is liable for breach of contract (conclusion).

How are judgements of the court—a collective body—to be fixed by judgements of its individual members? In other words, how are individual members' judgements to be *aggregated* into judgements of a single, collective entity (the court)? A natural mode of aggregation is majority vote. Suppose, then, that the three judges make individual judgements on the three propositions as shown in Table 1.1 below. Using the rule of majority vote, these individual judgements are then aggregated into collective judgements by the court as shown in the bottom row of the table. Here the collective agent is inconsistent. It affirms that the defendant was obliged not to act, affirms that he acted, but denies that the defendant is guilty. Given the legal doctrine, this set of judgements is inconsistent. However, each individual judge has a perfectly consistent set of judgements. Thus, there is inconsistency at the group level despite consistency at

	<i>Obliged not to act</i>	<i>Acted</i>	<i>Breached the contract</i>
<i>Judge 1</i>	True	True	True
<i>Judge 2</i>	True	False	False
<i>Judge 3</i>	False	True	False
<i>Court</i>	True	True	False

16

the individual level. Assuming, as is standard, that judgemental inconsistency is irrational, irrationality emerges at the group level.

The fly in the ointment, apparently, is the majority vote aggregation procedure. Surely there is another aggregation procedure—probably many of them—that won't create this problem! Unfortunately, this is not the case. There is no judgement aggregation procedure that averts the problem. List and Pettit (2002) first showed this with an impossibility theorem, and several related impossibility theorems have been proved as well (see Dietrich and List 2007, in particular). To illustrate, consider the following reasonable-looking requirements for a judgement aggregation procedure, which would take sets of individual judgements as inputs and generate a set of group judgements as outputs.

*Universal Domain.* The aggregation function takes as input any possible profile of attitudes towards the propositions on the agenda, assuming that individual attitudes are consistent and complete. [Completeness here means that for every proposition-negation pair on the agenda, the agent has a positive attitude on either the proposition or its negation.]

*Collective Rationality.* The aggregation function produces as output consistent and complete group attitudes towards propositions on the agenda.

*Anonymity.* All individuals' attitudes are given equal weight in determining the group attitudes. Formally, the aggregation function is invariant under permutations of any given profile of individual attitudes.

*Systematicity.* The group attitudes on each proposition depends only on individual attitudes towards it, not on their attitudes towards other propositions, and the same pattern of dependence holds for all propositions.

The surprising impossibility theorem is that no aggregation function satisfies all of these conditions (List and Pettit 2002, List and Pettit forthcoming).

As with other paradoxes in which initially plausible propositions or constraints are found to be jointly unsatisfiable, theorizing proceeds by trying to decide which propositions or constraints to abandon (see List and Pettit forthcoming). The most familiar example in recent epistemology is the sceptical paradox featuring deductive closure for knowledge



(Nozick 1981; DeRose 1995). Should the closure principle be rejected, or should another proposition be abandoned instead? In the present case we have a paradox of social epistemic rationality: Can there be systematically rational group agents whose judgements are suitably tied to those of their members? If not, what does this imply about epistemic rationality at the collective level? This question is sufficiently novel relative to epistemological tradition that it cannot be folded into preservationist SE. On the other hand, it is sufficiently continuous with traditional epistemology that it deserves a home somewhere in epistemology. A comfortable home can be found in what I call the expansionist district of SE.

## **5.2. Epistemic Evaluation of Social Systems**

Large chunks of analytic philosophy, epistemology included, focus on everyday thought and talk, often teasing out norms that implicitly govern them. Epistemology should not be restricted, however, to this activity. A broader view of epistemology is present in the epistemology of science, which definitely does not assume that lay practices are the exclusive or final benchmark of anything normative. Philosophers of science are more sympathetic to the *meliorative* project of improving epistemic practices, especially in the scientific arena.

In addition to science, there are other arenas where meliorative epistemological projects make eminently good sense. Many sectors of social life feature practices and institutions ostensibly dedicated to epistemic ends, but where one is entitled to wonder whether prevailing practices and institutions are optimal. Subjecting such practices and institutions to epistemic evaluation is therefore in order. Are they the best practices or institutions (of their type) that can be devised? What alternative practices would work better in epistemic terms? This kind of social epistemological project was advocated in *Knowledge in a Social World*—and might be just the sort of project that Alston's imagined critics would dismiss as 'not real' epistemology. Although I would grant that such projects are infrequently encountered in historical or mainstream epistemology, they would nonetheless be *continuous* with certain traditional epistemological projects. They would apply epistemic criteria of evaluation to admittedly non-standard kinds of objects: not individual agents or even collective doxastic agents, but social systems or policies that have a significant causal impact (for good or ill) on society's epistemic outcomes.

### **5.2.1. Legal Adjudication Systems**

Legal adjudication systems are an example of the kind of social system I have in mind. While not being collective doxastic agents in their own right, such systems are typically guided by a truth-seeking mission. In the case of the American system, the Supreme Court stated this clearly in *Tehan v. U.S.* (1966), 'The basic purpose of a trial is the determination of the truth.' How does the common-law trial system, used in Britain, the United States, and most former Commonwealth countries, go about trying to determine the truth? Which features of its design—at least according to its champions—make it

conducive to this end? And how well does the design actually work in promoting a high proportion of accurate verdicts?<sup>15</sup>

18

The crucial feature of the common-law system is its adversarial character. Under this system a trial is essentially a structured competition between opposing attorneys (or teams of attorneys) who represent the contending parties. Lawyers are responsible for all of the main initiatives and conduct of the case. They file briefs and motions, gather and prepare the evidence, and at trial their battle of interrogation and argumentation dominates the proceeding. The judge is a referee who enforces conformity with legal requirements by restricting introduction of evidence and constraining counsel's argumentation. The trier of fact, often a jury, remains essentially passive.

Proponents of the adversarial system claim that it's as good an institution as can be devised for achieving the epistemic goal of accuracy of judgements (verdicts). The best way for a court to discover the facts, according to the theory, is for each side to strive as hard as it can to present to the court all evidence that supports its claims. The neutral trier of fact then decides on the basis of the evidence presented. Attorneys are expected to be partisan representatives of their parties, but such partisanship is not inimical to the goal of truth. Dueling partisanship cancels out, so long as the trier of fact is neutral. Juror selection practices are designed to ensure juror neutrality.

The adversarial process partly utilizes the technique of dialectic or argumentation. This allows each side to try to persuade the trier of fact of its favoured conclusion(s) by adducing evidence for its thesis, rebutting the evidence of the opposition, and trying to defeat the defeaters offered by the opposition to its own evidence. What better way could there be for truth to emerge than to produce as much (relevant and reliable) evidence as possible and let the contending parties thrash out what it implies?<sup>16</sup> It is vital, of course, that all of the evidence be provided to the ultimate doxastic decision-maker. In a social fact-finding network, the network is likely to output an accurate verdict only if the network's designated decision-maker—in this case, the jury—receives and weighs all of that evidence. This is precisely how the system is designed to work. Jurors must sit through the entire proceeding and hear all of the evidence. They are not allowed to attend parts of the trial at their whim, choosing to attend only

19

when their favourite side (e.g. the prosecution) presents its case. They cannot cherry-pick the evidence, something that commonly happens in unstructured evidence-gathering contexts. In informal life people choose which newspapers to read, which talk shows to listen to, which friends to seek opinions from, and so forth, with no external enforcement of balance. The trial process aims to redress this defect in informal evidence gathering, and there are good epistemic grounds for this.

Thus, several features of the common-law system with its adversarial component promise epistemically good consequences. But meliorist social epistemologists should not stop

here. They should be alert to other features of the system that may detract from its positive epistemic prospects. Here are two problematic features.

The common-law jury is a lay jury. This contrasts with the composition of the fact-finder under the principal alternative tradition in Western legal systems: the so-called ‘inquisitorial’ or ‘civil law’ system found in much of continental Europe. The jury of the civil-law system traditionally consists of professional judges, who jointly occupy the three roles of investigator, judge, and jury. The difference between lay jurors and professional jurors is substantial. In complex cases—for example, product liability cases—there is typically a battle between scientific expert witnesses for the two sides, who give conflicting testimony on highly technical and arcane matters. Untrained and inexperienced jurors find it extremely difficult to assess the relative merits of the two sides of the argument. They are likely to perceive the battle as a deadlock; so they might ‘toss a coin’ to make a decision. If they do find one side more persuasive, will this be a result of appreciating genuine confirmational relations between evidence and hypotheses? Or will it merely result from rhetorical skill, polish, or impressive-sounding credentials? To generalize the problem, good epistemic upshots are not assured simply by giving doxastic decision-makers lots of evidence, even good-quality evidence. If they do not understand and interpret the import of the evidence correctly, the quality of the doxastic choices may still be poor. For this reason many countries do not use lay juries in complex tort cases.

A second possible weakness of the common-law system can be discerned from a closer look at its adversarial component. Just how good is this competitive procedure from an epistemic point of view? Are there defects hidden in the woodwork? Here is one defect. It is all well and good for both parties to have partisan counsel that serve their respective interests, but does this level the playing field so that truth is likely to emerge? This is the theory behind the adversarial system and what wins it many adherents; but there is a hole in the rationale. Legal representation is not of uniform quality. Some lawyers command higher prices than others, presumably because their skilled efforts produce a higher probability of winning a case than would the efforts of other lawyers. If so, disparities in the quality of legal representation can undermine the system's

20

reliability. This worry is commonly raised in connection with the poor quality of legal defence provided by public defenders. Whether through inexperience or excessive caseloads, it is generally thought that defendants represented by public defenders incur higher chances of conviction. This probably indicates that the truth may not be adequately served where there is a large disparity in legal resources utilized by the opponents.<sup>17</sup> Under the inquisitorial system, by contrast, where lawyers play a much less significant role in a trial, legal outcomes will be less sensitive to the quality of legal representation.

I don't mean to render an overall judgement about the comparative reliability of the adversarial and inquisitorial systems of legal adjudication. Obviously, these are highly

complex matters. The simple point being made is that there is work to be done in assessing the reliability properties of different legal adjudication systems. Now reliability assessments are a species of epistemological assessments.<sup>18</sup> A lot of epistemology's work is concerned with assessing psychological and formal 'objects' in terms of their truth-conduciveness. It's an epistemological question whether the process or faculty of 'intuition' is reliable. It's an epistemological question whether introspection is a genuine process of the mind, and if so whether it is reliable (and over what range of questions about the mind it is reliable). It's an epistemological question whether this or that inductive method, or statistical method, is reliable. To be sure, the project of evaluating social institutions and practices has not, generally speaking, been on the agenda of epistemology. That is why such projects should be assigned to an *expansionist* wing of social epistemology. It is clear, however, that it is naturally assigned to an expansionist wing of social *epistemology*, since the proposed mode of assessment is distinctively epistemological.

Can the quality of a legal adjudication system be assessed in exclusively epistemic terms? Many theorists would dispute this. Aren't considerations like fairness and impartiality crucial to the acceptability of a legal adjudication system? Yet surely fairness and impartiality are moral rather than epistemic factors. The critics would undoubtedly be right in saying that the criteria of a good legal adjudication system are not wholly epistemic. This doesn't show, however, that epistemic criteria are not highly significant, perhaps even paramount. Moreover, when examined closely, it is difficult to pinpoint the content of fairness or impartiality in dispute resolution without ultimately turning to epistemic desiderata (Goldman 1999: 280–1).<sup>19</sup>

### **5.2.2. An Epistemic Rationale for Freedom of Speech**

Expansionist SE of the 'systems' variety need not be restricted to the evaluation of entire social systems or institutions such as the common-law system of legal adjudication. It can focus equally on selected policy issues found in a social system. Consider, for example, a government's policy on freedom of speech. Just as many theorists have championed the adversarial system of legal dispute resolution by appeal to epistemic consequences, so it has been argued by both historical and contemporary writers that the diversity of ideas and debates that freedom of speech promotes provides the best way to test and determine their truth. Hence, freedom of speech is the best social arrangement for generating good 'veritistic' (true belief) consequences. One famous statement (roughly) to this effect was made in John Stuart Mill's *On Liberty*:

[T]he peculiar evil of silencing the expression of an opinion is, that it is robbing the human race; posterity as well as the existing generation; those who dissent from the opinion, still more than those who hold it. If the opinion is right, they are deprived of the opportunity of exchanging error for truth; if wrong, they lose, what is almost as great a benefit, the clearer perception and livelier impression of truth, produced by its collision with error. (1962: 142–3)

Similarly, Justice Oliver Wendell Holmes (in a dissenting opinion for the Supreme Court) wrote:

[W]hen men have realized that time has upset many fighting faiths, they may come to believe even more than they believe the very foundations of their own conduct that the ultimate good desired is better reached by free trade in ideas—that the best test of truth is the power of the thought to get itself accepted in the competition of the market. (*Abrams v. United States* 1919)

In other words, the best way for members of society to find and believe the truth is to have a free market for ideas in which speech goes unregulated (by the state). Both Mill and Holmes defend a certain type of state policy toward speech that is clearly couched in terms of epistemic consequences. It is therefore appropriate for social epistemologists to ask whether this epistemic-consequentialist claim is true—whether the deregulation of speech is the best way to maximize the acquisition of true belief (and error avoidance) in society. Although this could hardly be considered a well-recognized or well-trodden problem within

22

traditional epistemology, it is a perfectly good problem in expansionist social epistemology.<sup>20</sup>

### **5.2.3. Epistemic Approaches to Democracy**

As a final example of expansionist SE, consider the epistemic approach to democracy, a subject of active interest in contemporary democratic theory. A driving force behind an important strand of work on the epistemic approach is the Condorcet Jury Theorem, discovered in the eighteenth century. Suppose that members of a group are faced with a binary choice (a ‘yes/no’ question or a ‘true/false’ question) for which there is a correct answer. Condorcet proved that if the group takes a majority vote between the two alternatives, and if the voters all satisfy two conditions, *competence* and *independence*, then the probability that the majority picks the correct alternative increases to one (certainty) as the group size tends to infinity. What is meant by competence and independence? A voter is competent if she is more likely than chance to vote for the correct alternative. Independence is the condition that each vote is cast independently of the others. The Condorcet result is striking in its highly positive conclusion. It implies that large groups are virtually infallible even if their members are only slightly better than chance at selecting the better alternative. If majority judgements are indeed most probably correct in large societies, majoritarian democracy seems to receive strong support from an epistemic perspective.<sup>21</sup> In other words, democracy can be defended as a good political system not on the basis of some abstract principle of equal rights but because it does particularly well at making correct political decisions (e.g. decisions that best promote the society's interests).

Here is a way of spelling out the epistemic approach to democracy, one that makes the Condorcet Jury Theorem highly relevant. Jules Coleman and John Ferejohn (1986)

propose that democracy might be taken to have epistemic value in virtue of the fact (if it is a fact) that the outcome of a vote provides *evidence* for the proposition that the winning policy or alternative accords with (what Rousseau called) the ‘general will’ (or some other criterion of political rightness). If the conditions specified by the Condorcet Jury Theorem are met, it is plausible that the outcome of a vote does indeed provide evidence that the approved policy accords with the general will.

The usual knock against the Condorcet Jury Theorem is that it reveals that voting per se is not such a significant aspect of democratic procedure. The

23

crucial question is which binary choices are placed before the voters. This is characterized as an issue of *agenda setting*. How do political bodies decide which, of the infinitely many possible questions or choices to place before voters, are actually put on the ballot? Such decisions are generally thought to be a matter of debate and deliberation, in some forum or other, among some political actors or other. The serious work of a political body, then, resides in the conduct of public debate or deliberation: who should be involved in the deliberations and what rules or practices should govern the deliberative process (if the system is to be genuinely democratic). If public deliberation (of an approved kind) is a good epistemic process, an emphasis on deliberation might provide a satisfactory epistemic approach to democracy.

Many democratic theorists focus their attention on public deliberation, including Dewey (1927), Habermas (1996), Rawls (1993, 1997), Cohen (1989, 1997), Gutmann and Thompson (1996), Anderson (2006), Estlund (2008), Misak (2000), and Talisse (2005). Not all of these treatments of public deliberation invite the label ‘epistemic.’ In particular, those who focus on consensus as the aim or mark of good deliberation would not (thereby) earn this label, in my opinion. However, it might be possible to carve out a characterization of *reliable* (i.e. truth-conducive) public deliberation, realizations of which would fulfill the reliability desideratum emphasized by some proponents of the epistemic rationale for democracy (e.g. those inspired by the Condorcet Jury Theorem). I tentatively call this a *reliabilist approach* to democracy, and defenders of this kind of rationale or justification for democracy *reliability democrats*.

I shall briefly sketch the idea I have in mind, although there is neither time nor space to develop it fully. We can imagine groups of people who satisfy certain constraints on responsive democratic deliberation but do not thereby achieve the sought-after ‘epistemic goods,’ that is, reliability or increased reliability. For example, consider a group of ill-informed—indeed misinformed—citizens who set out to argue with one another on political matters and arrive at some sort of (rough) consensus. Everyone is allowed to speak, they listen carefully to one another, and even make epistemically justified doxastic responses to what they hear. Unfortunately, their misguided initial belief states are so bad—either because of a dearth of evidence, or an abundance of misleading evidence—that the deliberative process does not boost their average reliability with respect to the questions discussed. If this is possible (and I think it is), it shows that rational public

deliberation is not necessarily sufficient to yield reliable doxastic outputs, or even to increase reliability from a pre-deliberative state. However, it might be possible to add additional constraints on the standards of public deliberation so that if a group satisfied those constraints in addition to the first ones, public deliberation would tend to increase the group's reliability. Such a set of requirements for democratic deliberation would constitute a reliabilist standard of deliberation.

24

If a theorist managed to formulate such a set of requirements and showed that satisfying these requirements has reliabilistically beneficial properties for groups, this would be an *epistemological* contribution to democratic theory. It would show, by means of epistemology—indeed, social epistemology—that certain features of a system promote an important epistemic desideratum, that is, reliability, or increased reliability. The Condorcet Jury Theorem makes exactly this kind of contribution, although it does not address public deliberation or debate. Writers who emphasize the way that genuine democracy makes use of ‘situated knowledge’ to improve the community's overall knowledgeableability make roughly this kind of social-epistemological contribution to democratic theory. In other words, they may be reliability democrats in the sense specified above. Elizabeth Anderson (2006) provides a particularly good example of this.<sup>22</sup>

Is this kind of contribution to democratic theory a genuine (‘real’) piece of epistemology? It may be helpful here to distinguish between epistemological *theory* and *applied* epistemology, in parallel with the familiar distinction between ethical theory and applied ethics. Showing how some features of public deliberation would contribute to a group's reliability, thereby providing a basis for a reliabilist rationale for democracy, would be, at a minimum, a piece of *applied* epistemology.<sup>23</sup> The question of how to taxonomize such work is probably not pressing. But it is important to recognize the general category of applied social epistemology, and to be prepared to recognize work in that category as specimens of *real* epistemology.

## 6. REVISITING THE CRITICS: ‘IT AIN’T REAL EPISTEMOLOGY’

One would expect Alston's imagined critics to challenge the foregoing examples of expansionist SE. The refrain, again, would be, ‘This isn't the real McCoy; it isn't real epistemology’. However, it is not enough to say that these problems are not traditional parts of epistemology; this much has already been conceded. The question is whether these problems are sufficiently *continuous* with traditional epistemology. What would be the critics' specific bones of contention? Which of the examples adduced above would they dispute, and why?

An obvious and expected line of criticism would concern the choice of ‘targets’. There are only two suitable targets in real epistemology, critics might claim: (1)

25

beliefs and other psychological states, and (2) persons who are subjects of these states. Collective or group entities, the critics would continue, are not genuine subjects of psychological states, including doxastic states. So collective agents are not kosher targets for epistemological analysis or assessment. Finally, common-law systems of legal adjudication, free speech policies, and democratic systems are not psychological states or doxastic agents. So none of them is susceptible to being analysed or evaluated in epistemological terms.

What warrants the critics in restricting epistemological analysis or assessment to the two indicated types of targets? One obvious omission is (belief-forming) methods, a classical target of epistemic appraisal. True, collective entities and social systems are not methods. But social systems and institutions are something like methods insofar as methods are commonly evaluated by their (epistemic) upshots or outputs, which is exactly what is proposed for the epistemic evaluation of social systems. It is not claimed, of course, that traditional or mainstream epistemology has a longstanding preoccupation with evaluating social systems. But here we are discussing expansionist SE, which is presumed to mark a non-trivial departure from traditional epistemology. What is required for expansionist SE is that its modus operandi be significantly continuous with traditional epistemology. I argue that a reasonable continuity holds between methods and systems, each of which can be evaluated in terms of similar epistemic outputs.

Turn now to collective agents. This is not the occasion to launch a detailed metaphysical defence of collective doxastic agents. But if critics dispute the existence of such entities on ontological grounds, they will be picking a fight with a large community of philosophers. Moreover, this community of philosophers makes a strong case for such entities as the subjects of intentional attitudes. If a good case has been made for plural bearers of *certain* intentional attitudes, such as intentions, what are the principled grounds for denying the existence of plural bearers of doxastic attitudes?

What further objections might the critics have up their sleeves? Recall Alston's speculation that the imagined critics would 'relegate much of this so-called social epistemology to sociology, social psychology, or other social sciences, or perhaps to the philosophical foundations thereof'. Perhaps this would be their next tack. The sample problems I list under the heading of expansionist SE really belong to different disciplines entirely, not to philosophy.

Which disciplines do they have in mind? The problem about the adversarial system of legal adjudication would be assigned to law. The thesis about the epistemic upshots of freedom of speech might be assigned to economics, since proponents of the thesis often claim that it is derivable from economic theory. The epistemic approach to democracy would belong to political theory. The rationality of collective doxastic agents might be assigned to some mix of political theory, legal theory, and economics (especially social choice theory),



since the interesting paradoxes in that field take their inspiration from work in those disciplines. But aren't these disciplines entirely separate from philosophy? Wouldn't the critics therefore be right to say that at least these portions of SE belong outside of philosophy and hence outside of epistemology?

No. I accept the portion of the foregoing claim that says that the specified subjects belong *partly* to the specified, extra-philosophical disciplines. But this does not entail that they are 'entirely separate' from philosophy; I argue that they belong to philosophy *as well*. Several pairs (or triples) of these disciplines *overlap* or *intersect* at these junctures. All of these fields have theoretical or formal wings that intersect with applied portions of theoretical and formal parts of philosophy. It is no coincidence that the principal developers of the theory of group agency, Philip Pettit and Christian List, double as philosophers and political theorists. There are many other sectors of philosophy in which contemporary philosophers plow much the same turf as researchers in nominally different subjects. Logicians are found in both mathematics and philosophy departments. Decision theorists and game theorists are found in psychology, economics, and philosophy. Theoretical physics is debated in joint seminars involving physicists and philosophers. Psychologists, neuroscientists, and philosophers may all give papers at a single conference, and critique one another's claims. Similarly for philosophers and linguists. A recent book on the criminal law was written by an (erstwhile) philosopher of science and reviewed in a legal journal by a lawyer-philosopher team.<sup>24</sup> In many of these interdisciplinary interfaces, certain individuals specialize on the theoretical side of the field while others specialize on the empirical or policy side. Everyone recognizes, however, that satisfactory treatments of the issues require insights and angles of multiple kinds.

This laundry list of interdisciplinary research efforts is intended to highlight the fact that it's misguided to try to establish bright boundary lines between disciplines. Such bright lines would obscure the fact that a great deal of important research occurs at intersections of nominally separate disciplines. It is arbitrary to insist that such intersections be assigned to exactly one discipline. The next question is: which branch of philosophy do the topics we have surveyed belong to? By far the most plausible candidate, in each case, is epistemology. In most of the cases, a class of social objects is selected as targets of assessment, and in each case these objects are assessed (or guidelines are laid down for their assessment) in terms of characteristically epistemic criteria, such criteria as conduciveness to true belief or rationality of belief. In one case, the prospects are examined for justifying a certain type of political system (democracy) in explicitly epistemic terms. What other branch of philosophy do these investigations belong to if

27

not epistemology? Their credentials for being *real* epistemology are the use of characteristically epistemic criteria of evaluation or normative appraisal in all of the examples.<sup>25</sup>

28

## **2 Testimony, Advocacy, Ignorance: Thinking Ecologically about Social Knowledge**

Lorraine Code

I could not but wonder . . . what it would take to make my experience verifiable. The testimony of an independent white bystander? . . . The blind application of principles of neutrality. . . acted either to make me look crazy or to make the reader participate in old habits of cultural bias.

(Patricia J. Williams 1991: 47–8)<sup>1</sup>

### 1. POINTS OF ENTRY

Much early twenty-first century Anglo-American epistemology is residually loyal to a tradition for which perception, memory, reason, and testimony are the principal, most reliable sources of knowledge, and the epistemically autonomous, self-reliant ‘individual’ is the presumptive knowing subject. Testimony’s relegation to last place on this list signals its lesser ranking in relation to perception, memory, and reason (or induction), because reliance on testimony has long been thought to compromise epistemic self-reliance. Social epistemology unsettles this ranking. In granting a central place to testimony in knowledge production and circulation, it interrogates assumptions about the homogeneity and universal replicability of knowers and ‘the known’, while contending that epistemic self-reliance is at best an impossible dream, and contesting charges that testimony amounts to nothing more substantial than hearsay, rumour, or gossip. Persistent distrust of testimony is held in place by a tacit picture of the solitary epistemic agent whose self-sufficiency and open receptivity ensure that her or his direct perceptions in ‘normal’ or ideal observation conditions will yield reliable knowledge. Memory enables her or him to sort and match perceptions, which vary in reliability with the vividness of the originating sensory experiences. Yet testimony, even in a (Russellian) restricted form of ‘knowledge by description’,

29

is a risky affair even when it takes the form of corroborated oral or written reports.

This ingrained distrust of testimony is curious in view of how small a proportion of what people reasonably claim to know, ‘individually’ or communally, comes from or is verifiable in appeals to first-hand experience. It is more puzzling still when enquiry moves into the social register, away from the rarefied places where epistemologists commonly work and the abstract, sanitized examples on which they tend to base their conclusions: when it moves into places where testimony, in a richer than merely information-reporting sense, becomes the stuff of which knowledge is made and on which policies, practices, attitudes, and actions rely. There, ‘the knower’ cannot be represented as an isolated, self-sufficient being for testimony is, inherently, a form of address: it is explicitly and intentionally interactive. Hence it is a place where epistemic

individualism cannot hold. Although formal perception- and memory-stating propositional knowledge claims of the ‘*S*-knows-that-*p*’ variety may (implausibly) pose as monologic, as though no interlocutor(s) were required for their completion, a testimonial communication looks to its hearer(s) for acknowledgement, contestation, and uptake.

Social epistemology claims its title, in large measure, from the centrality it accords to testimony and to knowledge-conveying exchanges in the real world. It is striking for the attention many social epistemologists give to extended, situated examples of epistemic negotiation and for the subtle, far-reaching effects of linguistic shifts from impersonal, third-person propositional claims that ‘*S* knows that *p*’ to the language of speakers and hearers evident, to cite just two examples, in Edward Craig’s (1990) *Knowledge and the State of Nature*, and Miranda Fricker’s (2007) *Epistemic Injustice*. Where the *S*-knows-that-*p* rubric enabled formal early twentieth-century Anglo-American epistemology to transcend the vicissitudes of the world in order to discern a priori necessary and sufficient conditions for ‘knowledge in general’, social epistemologists return to and reclaim the world, both human and other-than-human, with its incoherence and messiness, its contradictions and specificities, to engage constructively and critically, descriptively and normatively, with real epistemic interactions and negotiations. Hence the very idea of ‘knowledge in general’ is drained of content. Social epistemology generates a range of issues that, for traditional epistemologists, counted merely as *hors de question*, many of which blur the dividing lines that have separated epistemological enquiry from ethical and political debate and influences. In social epistemology, ethical-political questions—such putatively extra-epistemological issues as trust, power, advocacy, negotiation, and epistemic community—inevitably enter the discourse, and not to the detriment of responsibly objective enquiry. Giving and receiving testimony commonly, if tacitly, involves many of these issues.

In his essay ‘Getting Told and Being Believed’, Richard Moran argues persuasively in favour of its being sufficient ‘for bringing someone to know that

*P* that they were *told* by someone who knew, and they believed him . . . if this second person is taken to *know* that *P*, he may tell another person, and so on’ (Moran 2006: 273).<sup>2</sup> Moran’s interest is less in conditions for the possibility of knowledge than in the *relation* of believing, ‘where its direct object is not a proposition but a person’ (Moran 2006: 273). His broader interest is in the relationships, the modalities of epistemic dependence and, I would add, of epistemic vulnerability and risk, that follow from reliance on testimony as a source of knowledge. In his view, ‘the distinctive relation of believing another person’ (Moran 2006: 274) is pivotal. This relation is quintessentially social, if sometimes in a minimal sense when relatively trivial, punctiform informational items are at issue. It is forged in communicative discourse, and hence in eschewing the *S*-knows-that-*p* rubric where any knower, interchangeably, can occupy the *S*-place and any object, the-*p* place. A certain particularity enters the scene, implying that, to a greater or

lesser extent according to the substance of the situation, it matters who that ‘other person’ is, for a responsible knower will not believe just anyone, unless for quotidian information—about the time of the train, the location of the bank. In invoking a ‘responsible knower’, however, I am not suggesting that the responsibility is the hearer’s alone. Equally urgent, and more significant for Moran, is the responsibility of the speaker in exchanges both trivial and complex: someone who tells me the meeting will take place tonight is giving me her *assurance*, he maintains, that it will happen. My informant participates as ‘a kind of guarantor, provid[ing] me with a characteristic reason to believe, different in kind from anything provided by evidence alone’ (Moran 2006: 279). Here, ethical issues are integral to the appeals to assurance and responsibility invoked in even the smallest moments of communicative action: introducing questions about assurance and responsibility underscores the extent to which analyses of testimony require a departure from epistemic individualism. Not only are giving assurance and taking responsibility interactive and communal, but their centrality to testimonial practices confirms that once the discourse of communicative engagement displaces the discourse of autonomy and rights, the very idea of epistemic self-reliance is opened to contestation. An epistemic stance for which responsibility is an overarching epistemic virtue<sup>3</sup> distinguishes itself, also, from reliabilism, which remains committed to individualism in conceiving of reliability as a function of innate intellectual capacities, and cognitive processes as quasi-mechanistic. On such a model, questions about responsibility need not arise.<sup>4</sup>

Implicit in the assumptions that govern Moran’s and Welbourne’s positions is a specific (if not specified) conception of ‘the social’ for which participants in testimonial knowledge-conveying and -receiving exchanges are situated on

a level playing field where these exchanges readily ‘go through’, and questions about perceptual/experiential commensurability are unlikely to arise. There is an established tradition in Anglo-American epistemology of beginning from the small and working toward the more complex, from a conviction that ‘complex’ amounts, in effect, to simple multiples of the small: to the small writ large, so to speak. There is a certain plausibility to such an approach, yet it is misleading in that it represents ‘the social’ as less than fully social, with none of the complexities of sociality. In glossing over the interactive character of communication and failing to consider how it is often engaged across intransigent social-conceptual differences where investigation and negotiation are required for understanding to be achieved, it leaves open the question just how ‘social’ its operative conception of the social really is.<sup>5</sup> It generates assumptions about the ease of communication and understanding—of uptake and acknowledgement (in Wittgenstein’s sense<sup>6</sup>)—fostering a conviction that ‘we’ are so alike that our knowledge conveying (in Welbourne’s sense<sup>7</sup>) is a process of straightforward transmission and reception on the model of a radio or television broadcast, where the speaker is a neutral medium through which knowledge is transmitted, packaged, perhaps, and thus touched by her particular rational (even creative) capacities, but untainted by her other idiosyncracies, biases,

predilections. The hearer is a passive, equivalently neutral receiver. Often, this state of affairs prevails, but more rarely, I suggest, than Moran or Welbourne allows. Hence it is unclear how they could address Patricia Williams's situation as she captures it in the epigraph to this chapter.

Williams's question is prompted by, and exemplary for understanding, a common public failure of acknowledgement and uptake, in this instance enacted in response to her testimony about an incident at a New York Benetton's store, where a buzzer system allowed sales personnel to refuse admission to 'undesirables'. Intending to make a purchase, she pressed her 'round brown face to the window' and buzzed for admission. A white, bubble-gum blowing teenager glanced her way, and mouthed 'We're closed', even though several white people were shopping in the store. The question she poses was prompted by the editors of a law review symposium on 'Excluded Voices' who, after making unacceptable (to Williams) cuts to her written analysis of the episode, indicated they did not publish 'things that were unverifiable'. In the final version of the essay—without consulting her—they sought to eliminate all references to race, contending that naming it was 'just a matter of style'. As a black professional woman in mid-1980s America, Williams was well positioned to recognize how racism was germane to such an incident, and acutely aware of how often charges of paranoia are levelled against racially marked people who allege race as a determining factor in

32

their treatment. Since her 'testimony' was to appear in a law review article, as an established professor in a prestigious US law school whose credentials, by default, become part of the tacit 'assurance' she offers of the accuracy of her testimony, she might without invoking these qualifications have anticipated a presumption in favour of her credibility. Clearly, in this case, believing a person is at issue, and factors integral to the sedimented 'social imaginary' of a racist society override any expectation that telling and hearing can suffice for testimony to establish its veracity.

In short, on Moran's criteria, Williams is disbelieved as a person. She may have been mistaken in the substance of her testimony, and additional evidence might show this to be so. She may have misread the Benetton situation, but this is not the issue for no attempt is made to establish or discredit such a possibility; and in this situation as it played out, the issue is less about logical possibilities of seeing otherwise than about the practical possibilities that influenced how her knowledge claims acquired public standing. 'Standards' are in place before the fact to determine her credibility, and her perception is that her race has overridden her professional status (which is presumably unknown to the teenager, but not to her law review peers). It also might be thought, in the initial encounter, to have overridden her class, which must have been legible from her appearance and could have told in favour of her trustworthiness. These are details people routinely take into account in information-seeking and testimony-evaluating activities.

When the editors insist that race should not be mentioned, in the interest of maintaining an (illusory) neutrality, it is impossible to evaluate Williams's testimony on criteria that

address the specificities of the social climate: specificities integral to the very possibility of such an event's taking place. These are public matters of credibility, incredulity, and the politics of belief: they are missing from Moran's and Welbourne's accounts owing to the kinds of example they rely on and the social imaginary within which their analyses are conducted. The larger issue, then, is about how social situations and 'identities' influence processes through which testimony can yield knowledge good of its kind. 'Principles of neutrality' that prohibit taking specificities into account render Williams's protests at best unreasonable, at worst—as she puts it—'crazy'. Hence they prevent this event and others like it from claiming a testimonial hearing on their own terms: no assurance she offers can break the epistemic deadlock.

Miranda Fricker's (2007) work on epistemic injustice develops a rich conceptual resource for thinking through such cases. Her analysis is grounded in a reading of virtue epistemology and an engagement with the implications of prejudice, incredulity, stereotypes, situated knowledge, and ethical-epistemological issues of trust, testimony, and social conditions for establishing the status of a credible knower. Epistemic injustice, in her view, has two principal modalities: testimonial and hermeneutical, each of which is a source of harm to subjects as knowers or as knowable. Contending that testimonial injustice necessarily

### 33

involves prejudice 'with the central case involving identity prejudice' (2007: 41), Fricker suggests that when the 'driving prejudicial stereotype involves the idea that the social type in question is humanly lesser [a woman, a black, a member of a marginalized class or group], the dimension of degradation *qua* human being is not simply symbolic; rather it is a literal part of the core epistemic insult' (2007: 44–5). She concludes: 'Testimonial injustice denies one access to what originally furnishes status as a knower. . . [it] can carry a symbolic weight to the effect that the speaker is less than a full epistemic subject' (2007: 145), incapable, then, of offering the tacit assurance Moran describes, that renders a testifier appropriately believable as a person.<sup>8</sup>

Hermeneutical injustice is similar to and different from testimonial injustice. It is, I suggest, also operative in the Williams example. It centres on the claim that 'relations of power can constrain women's [and members of other marginalized social groups'] ability to understand their own experience' (Fricker 2007: 147). (Williams's testimonial evidence suggests that she understands her own experience well, but is hermeneutically blocked in attempting to persuade her interlocutors of the veracity of her understanding. Hermeneutical injustice functions variously at different social layers.) For Fricker, testimonial and hermeneutical modalities of epistemic injustice have an 'identity constructive power in common', but they are structurally distinct in that testimonial injustice is 'inflicted individual to individual', while hermeneutical injustice 'is not inflicted by any agent, but . . . is caused by a feature of the collective hermeneutical resource' (Fricker 2007: 168). It draws on socially embedded interpretations and

understandings which are differentially available across relations of power and privilege. Thus Fricker maintains:

relations of unequal power can skew shared hermeneutical resources so that the powerful tend to have appropriate understandings of their experience ready to draw on as they make sense of their social experiences, whereas the powerless are more likely to find themselves having some social experiences through a glass darkly, with at best ill-fitting meanings to draw on in the effort to render them intelligible. (Fricker 2007: 148)

As Williams's story shows, power and powerlessness are fluctuating signifiers: as a relatively powerful member of her profession, Williams can make sense of her social experiences; but in certain public situations and venues she is relegated all over again to the powerlessness of her black American history.

According to Fricker, a collective social imagination functions as a repository for stereotypes, some benign or neutral, others unequivocally malign, and many of which 'shape our credibility judgements by stealth' (2007: 38–9), unknowingly. Thus, hermeneutical injustice can be systematic or incidental, but either way it inflicts harm to a person's or a group's very sense of self. While agreeing with her in the main, I find greater explanatory resources in the idea of a social *imaginary* to which I alluded earlier: a conceptual-analytic resource analogous but not identical

34

to the social imagination. As it is enlisted by Cornelius Castoriadis, to whom I am indebted, a social imaginary refers to implicit but effective systems of images, meanings, metaphors, and interlocking explanations-expectations woven through a social-political order, within which people in specific time periods and cultural-geographical situations enact their knowledge and subjectivities, and craft their self-understandings.<sup>9</sup> An *instituted* imaginary carries the normative social meanings, customs, expectations, assumptions, values, prohibitions, and permissions—the habitus and ethos—into which people are nurtured from childhood, which they affirm, challenge, internalize, or reject as they make sense of their place, responsibilities, and options, within a social and physical world whose 'nature' and meaning are also instituted in these imaginary significations. By contrast, an *instituting* imaginary is a locus of social critique and change: it refers to the critical-creative activities of a society recognizing that it is incongruous with itself, with scant reason for self-satisfaction, exhibiting its autonomy in exercising its capacity to put itself in question. Imaginatively initiated counter-possibilities interrogate the social structure to destabilize its pretensions to naturalness and wholeness, to initiate a new making. In my view, a social imaginary carries greater explanatory and revisionary power than the social imagination: it is more thoroughly systemic in a quasi-Foucauldian sense, more explicitly power-infused and politically oriented, less presumptively cerebral, and less conceptually contained than the idea of the social imagination. Where Fricker holds ethics and politics apart, I think of them as co-constitutive in ways that the 'social imaginary' captures.

## 2. ECOLOGICAL THINKING

Diverse moves toward naturalizing epistemology that originate with W. V. O Quine (1969b, 1994) have prepared the ground for the position I articulate in *Ecological Thinking*<sup>10</sup> where I am indebted to and critical of Quinean naturalism. In a pathbreaking departure from the reductive abstraction and formal analysis that have held orthodox Anglo-American epistemology aloof from the very knowledge it aims to explicate, naturalized epistemology shifts the focus on to questions about the nature and scope of real knowing. Constrained neither to analysing what ideal knowers ought to do nor to silencing the sceptic, naturalists assume knowledge is possible and seek to explicate and evaluate its real-world ('natural') conditions. Hence some feminist and other postcolonial theorists, and I among them, disenchanted with orthodox epistemology's incapacity to engage with 'situated knowledges'<sup>11</sup> in the epistemologically and socially-politically pertinent dimensions of their situatedness, have made common cause with them.

35

Many of these philosophers see social epistemology as a version of naturalized epistemology, which aims to investigate the constitutive part specific social circumstances play in the production and evaluation of knowledge: to examine how they influence or even determine who can participate in epistemic projects and who can claim a hearing—can expect acknowledgement—in testimonial exchanges. Social epistemologists investigate the structures and implications of the political, economic, and social settings where enquiry is conducted, engaging critically with the ideological commitments and the availability of models and narrative forms in the instituted imaginary that shape and influence the structure and content of observation and enquiry. They examine the socially conferred authority of diverse modes of investigation and styles of reasoning and the patterns of scepticism and silencing that govern knowledge construction. Analogously, and often practised as both socialized and naturalized, feminist epistemology examines the effects of socially produced gender norms and gender-specific interests and experiences in the production of knowledge.<sup>12</sup> Such investigations are continuous with inquiries into the effects of socially constructed ways of knowing or ignoring race, class, ethnicity, and other Otherings from the neutral, faceless, and unidentified knower who occupies the *S* place in standard epistemology. As I argue in *Ecological Thinking*, they encompass investigations of the effects of place, of situation, for the production and circulation of knowledge. Many such projects work with elaborated case studies and proceed by analogical argument from case to case, contending that standard attenuated 'one-liner' examples are too limited in their explanatory capacity to yield socially responsible knowledge and understanding.

Ecological thinking works across a range of subject matters and disciplinary territories, often acting as a scavenger in its quest for viable epistemic resources. It requires knowing the detail of place, population, politics, and particularity as enabling or constraining factors in knowledge-gathering projects. It emerges from and addresses multiple interwoven, sometimes contradictory social-epistemological positionings—feminist,



postcolonial, racist, classist, environmental, sexist—with the result that its detail and implications require multifaceted chartings. Yet ecological thinking is not just *about* ecology or *about* ‘the environment’. It generates a revisioned engagement with knowledge, subjectivity, politics, ethics, science, citizenship, and agency that pervades and reconfigures theory and practice in multiple, intersecting projects. It does not reduce to a set of rules or methods; it may play out differently from location to location; but its operative principles are sufficiently clear, and it is sufficiently coherent to be interpreted, enacted, and evaluated across widely diverse situations. It engages with events and situations across the human and other-than-human world, in projects where the dividing line between the (usually quantitative) *Naturwissenschaften* (natural sciences) and the (commonly

36

interpretive) *Geisteswissenschaften* (human sciences) blurs, and epistemic and ethical-political matters are reciprocally informative. In such projects, justification is more often pragmatic than deductive-nomological, although not opportunistically so, and governing principles are akin to those that govern virtue epistemology and ethics, enacted in media res and by example, rather than on normative a priori principles. Hence, the evaluative structure it yields is something of a hybrid.

Within this conceptual frame, I am interested in how *doing epistemic justice* to people silenced, ignored, or subjugated by received conceptions of knowledge, and to circumstances and places that fall outside the purview of what is readily knowable in spectator epistemologies, is complicated by incredulity, ignorance, and mistrust. Often, a curious *politics of unknowing* silently contributes to how people recognize and respond to ‘difference’: to judgements about whose knowledge matters or deserves a hearing, whose putative knowledge warrants or fails to claim acknowledgement, in intransigent social structures of scepticism and disbelief. It produces an epistemic inertia, a resistance to looking beyond the instituted imaginings that hold received views in place and to taking responsibility for their persistence. The Williams example is pertinent in this regard.

Although I call my position *ecological naturalism*, and situate it in the vicinity of Quine's epistemology naturalized and its successors, I resist Quineans' adherence to a scientific model of knowledge and the ‘unnaturalness’ of their conceptions of ‘the natural’, rather as I resist the ‘a-social’ conceptions of the ‘social’ operative in Moran's and Welbourne's analyses.<sup>13</sup> Quinean naturalism is inhospitable to the social-political critique of institutions of knowledge production and knowledge-producing practices which is integral, in my view, to the very idea of social epistemology, and which figures centrally in ecological thinking. Indebted methodologically and ideologically to the science of ecology, environmentalism, and the ethical-political impetus that inspired the new social movements in the second half of the twentieth century, ecological naturalism offers more ‘natural’ analyses of human epistemic practices and their products than formal epistemology or Quinean naturalism can offer. Taking as given that human beings are naturally, necessarily social beings, it proposes richer, less reductive possibilities for

socially responsible, transformative epistemology than are available in theories of knowledge modelled on an idealized scientific method.

Eschewing the residual positivism whose traces linger in Quinean naturalism, for which interchangeable observers leave their subjectivity and agency, interests and enthusiasms, actions and interactions, prejudices and hopes outside the laboratory door, to isolate their disinterested epistemic practices from what Bruno Latour (2004) aptly calls ‘matters of concern’,<sup>14</sup> ecological naturalism

37

looks to the new (contrasted with physics, chemistry, geology, astronomy) science of ecology. It is politically engaged, accountable for the knowledge it produces, often explicitly and unabashedly conducting its enquiries in the service of values, commitments, agendas, and political programmes that have themselves to be kept open to critical, deliberative evaluation. Its appeals to case studies where knowers are, by definition, precisely situated as active participants in producing and testing knowledge claims makes it implausible to imagine them isolated from the knowledge they produce, even when they appeal to laboratory-derived and other forms of empirical/experimental evidence for pieces of it. In these and other respects, ecological thinking is continuous with and makes common cause with feminist theories of strong objectivity<sup>15</sup> and situated knowledges.

Rachel Carson's epistemic practice is exemplary for ecological thinking. Her method of working back and forth from field studies to scientific analyses, experiments, and explanations, to specifically located experiential (testimonial) evidence, and local (or larger) histories of the sites and species she studies shows what I mean. The place she accords to secular testimony is germane to this point. Carson's scientific world was structured by hierarchical evaluative divisions between fact and anecdote, truth and narrative, reason and feeling, with the first of each pair claiming a level of credibility and authority to which the second could not aspire. Yet in her respect for testimonial—experiential—evidence of the everyday, down-on-the-ground variety commonly relegated to the merely anecdotal, epistemologically irrelevant, Carson contests and unsettles such distinctions (and earns the epithet ‘subversive’ in so doing).<sup>16</sup> Denigrations of testimony as evidence derive, if curiously, from the (lowly) status of experience in putatively experience-based epistemologies which accord minimal evidential respect to individuated, first-person accounts as they contrast with formal experience that is everyone's and no one's. Its formality protects it from the confusion and bias of particularity. Patricia Williams's experience, enacted within structures of power and privilege which valorize *some* people's testimony and denigrate others' in the name of an illusory neutrality, exemplifies this pattern, posing vexed questions about how experience can stand as evidence in an epistemic climate where credibility is unevenly distributed and testimony (again, *contra* Moran) is as often discounted because of whose it is as for its content. The issue translates, with variations, into current feminist, anti-racist, and other postcolonial theories of knowledge, and into the examples I will now discuss.

In Carson's scientific practice, as in biologist Karen Messing's research into women's workplace health and in the epidemiological field work of a Canadian IDRC team in Tanzania,<sup>17</sup> local testimony and orthodox scientific enquiry confront one another in ways that are instructive for understanding the place

38

of testimony in epistemic situations more elaborated, and more public, than one-on-one exchanges or unidirectional transmissions. In each case, albeit differently in each, advocacy—which is still more epistemologically suspect than testimony—is integral to enabling testimonially derived knowledge claims to go through: to achieve acknowledgement.

In her research into the effects of post-Second World War pesticides, Carson concludes on the basis of wide-ranging local testimony that citizens 'often show a keener understanding of the dangers and inconsistencies of spraying than do the officials who order it done' (Carson 1962: 113). Her investigations of massive insect-control spraying expose an interplay between precise secular hypotheses and larger empirical-experimental conclusions. Detailing a near-fanatical public resistance to knowing the 'almost unparalleled wildlife destruction' (Carson 1962: 92) that the eradication programmes leave in their wake, Carson poses 'a question that is not only scientific but moral . . . [W]hether any civilization can wage relentless war on life without destroying itself, and without losing the right to be called civilized' (1962: 99). In *Silent Spring*, she cites scientific-observational evidence and '*anecdotal evidence*. . . from nonexperts—suburbanites and ornithologists upset by the dead and dying birds they found on their lawns. . . campers poisoned by DDT' (Raglon 1997: 206, my italics). ('Her conclusions were based not only on the work of scientists on the cutting edge of their fields but on the case histories that filled her file cabinets', her biographer notes [Lear 1977: 357].) For Carson, testimony can set enquiry in motion and function as an ongoing corrective: it generates precise local hypotheses, confirms or contests interim findings. Moving back and forth between testimony and science, then, makes her (interim) empirical conclusions possible. So, in *Silent Spring*, Carson recommends evaluating 'some of the major control programs and learn[ing] from observers familiar with the ways of wildlife, and unbiased in favor of chemicals, just what has happened in the wake of a rain of poison falling from the skies' (Carson 1962: 86). Presumably, standard epistemology has the capacity to address routine failures to rely on scientifically uninformed lay testimony, but Carson is working in a climate where such recognition rarely occurs, and it is with this down-on-the-ground reluctance that she must contend to achieve credibility for her findings.

Noteworthy is the part testimony plays in enabling Carson to negotiate a hearing for knowledge claims to which the scientific establishment in post-war United States was stubbornly resistant. Much of this testimony is indeed anecdotally conveyed by testifiers for whom the reported incidents (recalling Latour) are both matters of fact and matters of concern. Carson evidently works from a presumption that the observers who report to her are believable; that in testifying to the damage, they are offering the assurance of 'a

guarantor' providing her 'with a characteristic reason to believe, different in kind from anything provided by [scientific] evidence alone' (cf. Moran 2006: 279). The sheer accumulation of such reports, and their capacity to withstand empirical

39

scrutiny, lends credence to her presumption. But neither her credibility nor that of her informers goes uncontested: one challenger, who declared she must be a communist, is quoted as saying: 'As for insects, isn't it just like a woman to be scared to death of a few little bugs! As long as we have the H-bomb everything will be OK.'<sup>18</sup> I cite the remark to underscore the epistemic insult enacted by (gender) identity prejudice in discounting/dismissing knowledge claims by discrediting their maker (Fricker 2007: 41), and to point toward the complexity of epistemic negotiations required to establish the veracity of testimony that goes against the grain of received opinion, in its substance or in the persona of the testifier.

As Carson maps the local and wider effects of pesticides, Messing (1998) maps the physiological effects of workplace conditions in the lives of poultry processors, waitresses, bank tellers, telephone operators, and others. As Carson relies on scientific-observational and testimonial anecdotal evidence, so Messing and her colleagues rely on scientific-observational and experiential evidence culled from interviews with workers, union members, and other 'ordinary folk', to investigate how 'women workers and scientists can be brought together to make occupational health studies more accurate and effective in improving working conditions' (Messing 1998: xvi). When symptoms are specific to small groups of workers, they afford insufficient (statistical) certainty either to establish the 'reality' of suffering, or to support workers' compensation claims. Hence Messing investigates how minutely individuated complaints play out for specific people (usually women) in the particularities of workplaces, habitats, governed by an ethos of profit and efficiency, backed by an epistemology reliant on statistical certainty alone. Her practice, like Carson's, is as subversive as it is respectful of the norms of scientific enquiry and in the place it accords to testimony in knowledge production. Messing reads statistical evidence against the elusiveness of certain symptoms to inclusion under law-like descriptions, moving between the laboratory and the workplace floor as Carson moves between the laboratory and the field or the ocean floor. Her work, too, exhibits an interplay between local hypotheses and empirical generalizations, where analogies traced and tested move enquiry forward, disanalogies initiate productive rethinking. At issue, then, is an investigation of the traditional 'context of discovery', where questions about how statistical significance is established need to be addressed. Statistics, too, are a product of social judgement and deliberation: they are not self-announcing.

Frustrated by attempts to convince unions, employers, and workers themselves of the *factuality* of correlations between workplace conditions and women's health, Messing observes a vicious circle in which 'researchers do not think of looking for an occupational cause . . . such causes are not demonstrated, so the

40

problems are attributed to women's biological or psychological “nature” ’ (1998: 19). She aims to undermine the injustices such stereotype-reliant ‘naturalizing’ (1998: 25) enacts, to show that women's workplace health is less about generic female incapacities, more about standards of fitness based on fixed ideas about equipment, work environments and ‘attitudes’ derived from what is typical for men. Thus, she shows how work stations, tools, and equipment designed for the ‘average man’ place many women and some men at a disadvantage, whose effects in pain, stress, and discomfort are rarely acknowledged except in the language of failing to meet an abstract, one-size-fits-all criterion of ‘fitness’.

The politics are complex. Workers who turn to activism to protest their circumstances meet with suspicions that ‘they will fake symptoms to gain their point . . . even when the patterns of physiological change are so specific to the toxic effect that the worker would have to be a specialist . . . to produce them’ (Messing 1998: 69); a mail sorter reporting extreme arm, shoulder, and elbow pain loses a compensation claim because an expert witness ‘quoted from the scientific literature purporting to show that repeated movements could not produce injury below a specific weight manipulated . . . [and] described his own laboratory tests, which . . . showed that this type of work was not dangerous’ (Messing 1998: xv). No one examined the sorter's situation (1998: xvi); yet without the knowledge that only such an examination could provide, the reports claim dubious validity. Hence Messing asks: ‘*how do scientists decide whom and what to study in the occupational health field?*’ (1998: xv). Adequate answers require attending to the workers' testimony, with an initial presumption in favour of their believability; analysing their interactions and situatedness ecologically, from a research stance informed and vigilant to detect local, if minuscule, epistemic injustices—testimonial injustices—generated by the politics of knowledge (gendered, raced, classed, aged, ableist and other) and the profit-driven interests that preserve invisibility and insignificance, credulity and incredulity across populations and occupations.

In my third example, a 1990s IDRC project in Tanzania, local informants had been ignored and their testimony denied acknowledgement in a climate of pervasive, stereotype-derived scepticism. An epistemological reversal, initiated by the IDRC, turned the health system around, moving it from ossified methods of gathering, evaluating, and circulating knowledge and outmoded administrative practices and distributions of epistemic power and authority, toward a responsible democratic complex of social-natural epistemic interactions. Each stage required negotiations with intransigent bureaucratic administrations and efficiency-driven multinational funding agencies, en route to repairing lacunae in an intransigent hermeneutical repertoire.<sup>19</sup> Here I sketch the project's epistemic and moral-political trajectory.

By the 1990s, Tanzanians had endured a health crisis for most of a generation, even as HIV/AIDS and recurrent infectious diseases were ‘brutally rearrang[ing] the social landscape’. Nor was new, centrally administered funding the solution, for the problem, according to the IDRC ‘FIXING HEALTH SYSTEMS’ report,<sup>20</sup> was an entrenched

epidemiological theory, rooted in statistically driven central planning, implemented from afar. The putative knowledge informing this theory and the epistemic injustices it produced are instructive. The project unsettled the power of conventional wisdom and its practitioners in ways that are ecologically responsive to the detail of place and demography and naturalized in a quintessentially down-on-the-ground fashion. Yet its success can mask the intricate epistemic negotiations and advocacy without which it could not have come about. My aim is to make them visible.

The IDRC'S simple recommendation was that monetary investments should be based on evidence that targeted the local 'burden of disease...in a particular ecosystem': administering aid from above, without such local understanding, was a principal source of the crisis. Established epistemic practice in its relentless, quasi-automatic repetition and rarely contested hegemony, had installed a screen of unknowing and not needing to know between health care administrators and local practitioners, pathologies, and places.<sup>21</sup> Hence, new approaches to local circumstances, sensitive to their specificity, were required: understanding the peculiarities and interconnections among diseases, and learning to hear, interpret, and act upon testimonial evidence from sources commonly denied authoritative voice under the Western eyes of development agencies and their centralized bureaucracies. (As Fricker notes, testimonial injustice can be 'pre-emptive' when potential interlocutors are ignored; their testimony is not solicited. A 'speaker is silenced by the identity prejudice that undermines her credibility in advance' (Fricker 2007: 130).)

In Tanzania, fixed ways of presuming to know relations of identity and power had produced a situation where local people could not dissent from distorted understandings of their experiences and circumstances. In one-size-fits-all aid distribution practices, funds were paid into a central administrative structure, on the assumption that appropriate local allocations would follow. But 'appropriateness' meant despatching equal sums of money and identical packages of drugs to each district, without investigating the specificity of populations and

health issues, which varied markedly across this vast country. A presumption of human sameness and of the reliability of bureaucratic knowledge generated a reluctance to engage directly with people who might be thought to know their situations and the obstacles they faced, yet had no access to the discourse table, no status as knowledgeable testifiers. All clinics received the same drugs regardless of whether the district's patterns of disease and death showed those to be the right ones to maintain health and reduce mortality there. Ignorance of the 'fit' between local causes of mortality and resource allocations saw non-malarious highland areas receiving a full complement of anti-malarial drugs more suited to an endemic area. While malaria, responsible for 30 per cent of life-years lost, was receiving 5 per cent of one district's budget, childhood diseases such as diarrhoea, pneumonia, and malnutrition, responsible for 28 per cent of deaths, were receiving 13 per cent, and tuberculosis, at 4 per cent of the burden of disease, 22 per

cent of the funds. People were sick and dying, and health-care workers, unable to offer appropriate treatment, were demoralized. A sceptic might contest my proposal that a dysfunctional epistemology is at work in this case with the counter-claim that no epistemology would sanction such a system. But it is not so easy to gainsay the evidence, first, that in the Tanzanian situation such an epistemology indeed informed and sanctioned the operative system; and secondly, that the practical circumstances that influence how certain misguided knowledge claims acquire public standing are the issue here, for social epistemology. Logical possibilities have little to do with the case, and standard epistemology is incapable of addressing it.

The IDRC's simple proposal was to study separate localities in their particularity, to determine whether epidemiological detail varied sufficiently to discredit a one-size-fits-all approach. Selecting two contrasting rural districts—Rufiji, with a dry, flat interior and a tidal delta on its coastline, and Morogoro, which is mountainous and lush—researchers mapped diseases in their regionally specific manifestations. Yet, maintaining an impressive level of local sensitivity, their aim was not to ‘apply’ locally achieved knowledge to other localities, but to work by analogy from district to district, testing conclusions for their adaptability ‘given the appropriate local statistical inputs’ (8/14).

The ‘mapping model’, described as ‘evidence-based’, differs sharply from Evidence-Based Medicine (EBM) that has generated debate in the United States and elsewhere. While textbooks praise EBM for eschewing such ‘subjective’ practices as interpretation and narrative,<sup>22</sup> Tanzania had neither the statistical tools nor the scientific approach to do the investigations orthodox EBM requires, nor to perform Random Control Trials (RCTs). Moreover, even with such assumptions and tools, it is unclear that interpretation-free analysis could have filled the hermeneutical lacunae, for eighty per cent of Tanzanian deaths

occur at home, and many fail to make their way into official documents or reports.

The epistemological issue is not about correcting isolated factual errors, but determining how outworn templates could make room for ecosystem-specific evidence and ecologically responsive practices. The questions are as political and ethical as epistemological: the power/knowledge inequalities between NGOs and their beneficiaries reinforced the ‘screen of unknowing’ I have mentioned, preventing NGO personnel from seeing—indeed, excusing their not-seeing—what their would-be beneficiaries needed, in circumstances where their voices were rarely accorded the credibility owed to presumptively reliable testifiers. Developing a solution required engaging with local villagers and clinicians as believable informants, in respectful evidence-gathering negotiations whose deliverances were perhaps less objectively ‘accurate’ by first-world standards than statistical analysis is imagined to be, yet were capable of withstanding epistemic and practical-political scrutiny. The IDRC report confirms, ‘when communities are directly involved in identifying and solving their own problems . . . [their] members become a powerful force in programs of social improvement’.

As an exercise in naturalized-social epistemology with the respect it accords to testimony, the IDRC's proposal was ingenious. When a death occurred, trained local researchers travelled by bicycle, motorcycle, or boat to conduct 'verbal autopsies' with survivors. Because of the sensitivity required to interview bereaved families well, the autopsies might take several hours and require repeated visits. Yet their meticulous, respectful laboriousness is the source of their capacity to avoid or erase epistemic injustice. Following the autopsies, findings were reviewed by three independent physicians to produce an epidemiological picture by feeding the conclusions into a computer according to a 'standardized international format'. The investigators allow that they rely on 'fairly accurate guesses' such as EBM aims to avoid, yet follow-up analyses confirm that, in aggregate, verbal autopsies provide a remarkably reliable picture of disease in the population. Increasingly, such Demographic Surveillance Systems (DSS) are used in Tanzania and elsewhere, both to compile health information, and to monitor poverty levels, education, food security, and the environment.

An additional obstacle in this power/knowledge complex was a reliance on stereotypes embedded in the Western colonial and postcolonial social imaginary, through which administrators and other outsiders presume to know local populations. Such 'webs of belief' with their thinly veiled racist-colonialist tenor were informed and shaped by conventional (outsider) wisdom for which death rates in these districts attested to Tanzanians' 'stubborn preference for traditional healers over modern health care'. The assumption was reinforced by a claim that even more 'enlightened' Tanzanians who seek modern care for malaria associate convulsions of late-stage malarial fever—*degedege*—with evil spirits or changes in the weather. Hence, administrators had concentrated on

44

undermining local faith in traditional care, thereby compounding a virulent mix of epistemic injustice enacted in observing the population from a distance as spectator epistemologies do, drawing conclusions from actions arbitrarily selected to confirm a prejudice-infused hypothesis kept alive by a colonial history of attributing naive innocence to putatively pre-scientific people. It combines with the pre-emptive hermeneutical injustice enacted in failing to work with or listen to indigenous people as reliable informants.

The Demographic Surveillance System presents a different view. Testimonial evidence shows Tanzanians who sought modern health care greatly outnumbering those who do not. It vindicates the scepticism many researchers had voiced about reliance solely on clinical attendance and government-compiled cause of death statistics as the knowledge-base of health planning. When local informants testify to problems of access to health care facilities and delayed responses to requests for treatment, or to the sheer inability of modern facilities to prevent some patients from dying, the emphasis shifts to understanding local circumstances and practices in and where they are and providing accessible care, responsive to local specificity.



### 3. TESTIMONY, ADVOCACY, IGNORANCE

These examples highlight issues that acquire centrality once epistemology re-engages with the complexities of the world. Moreover, in each, if differently, advocacy of some kind counts among the conditions that enable knowledge claims to go through; and in each, patterns of ignorance tacitly preserve an entrenched and silently operative social imaginary, to generate a certain inertia in the face of imperatives to know better. In this section I address some of these issues as they pertain to these examples and to testimony and epistemic in/justice more generally.

In *Ecological Thinking* (2006) I propose that *practices of advocacy often make knowledge possible* in hierarchical distributions of epistemic autonomy and authority and institutional divisions of intellectual labour, in Western societies.<sup>23</sup> The claim is contentious, for advocacy is a tainted practice for epistemologists: it recalls Plato's quarrel with the Sophists, and is one which positivists and their heirs were committed to suppressing. Connotations derived from legal argumentation, where an advocate argues someone's or some group's cause with the aim of winning the case represent advocacy as a practice which sacrifices truth, factuality, to achieving that goal. Such contentions are often justified: they prevail in the public and the philosophical imagination. But in the extended meanings and modalities that inform my argument, advocacy moves into epistemic lives and situations where it has to do with enlisting accredited expertise, adducing

45

evidence to support silenced or discounted ways of understanding experiences and phenomena.<sup>24</sup> It contributes to knowing a cause, a policy, a point of view well; and often to speaking on behalf of, supporting, a project or policy; to representing someone/some group so as to counter silencing, incredulity, and other egregious epistemic, personal, and social harms. It may be enacted in individual and communal practices: someone may advocate on her/his own behalf or on behalf of (an)other person(s); advocate for or against the significance, validity, credibility of someone's testimony, or of the testimony of a group, an institution, or society. In this last range of practices, advocacy carries epistemological force in the power and privilege structures of hierarchical societies, where it can become a liberatory practice whose goal is to establish the credibility of, and hence to (re)enfranchise epistemically marginalized Others.

Although advocacy often succeeds in opening epistemic space for the disenfranchised to achieve a hearing and for challenging the limits of entrenched epistemologies of mastery, advocacy advocates must be ever mindful of its perils: of 'its liability to the perversions of interest group prejudices, coercive paternalism, manipulation and appropriation' (borrowing Peta Bowden's words),<sup>25</sup> and of how its reliance on trust makes of it a fragile process whose participants are peculiarly vulnerable (cf. 2006: 176ff). I retain my brief for advocacy while acknowledging the seriousness of these charges, mainly because the implied contrast is implausible—between epistemic projects where risk and vulnerability do not figure because knowledge claims readily 'go through', leaving their claimants

untouched by the knowing, and advocacy practices where they are risky ventures. Advocacy on both ‘sides’—for the advocate or the advocated-for—requires ongoing strategic scepticism, wariness;<sup>26</sup> but the same is true, if differently, of most knowledge producing practices beyond the abstract, formal structures of *S* knows that *p*; and especially of those involving Latour's ‘matters of concern’.

Advocacy is vital to the (often partial) knowledge that issues from the epistemic negotiations I have discussed. Patricia Williams had to advocate, not altogether successfully or without risk, for her testimony to reach the public eye, to be evaluated there. Rachel Carson's tireless efforts in advocating for recognition of, and action in response to, local testimony about damage enacted by indiscriminate use of pesticides were both successful in placing ‘new knowledge’ before the public eye, and personally damaging. Karen Messing's advocacy for the ‘truth’ of women's workplace testimony, and the IDRC's for the validity of local testimony, with the practical actions they went on to inform, further support my claim in its strongest sense: *practices of advocacy often make knowledge possible*. In short, without persistent advocacy, Williams might have failed to make her experience

46

known; without persistent advocacy in writing, and in the public political forum, Carson would not have achieved the acknowledgement (in Wittgenstein's sense) from the Kennedy government that resulted in a ban on pesticides; Messing would not have convinced public granting bodies to recognize the factual import of workers' ‘statistically insignificant’ testimonial reports; and the IDRC would not have succeeded in bringing local Tanzanians to testify at the discourse table.

‘Advocacy practices’, then—representing, arguing for, recommending, acting or engaging in projects of enquiry and intervention with, in support of, or on behalf of some one or some group—animate an *instituting* social imaginary capable of unsettling the inertia of an entrenched imaginary, thence promoting social conditions of co-habitability in the best sense, upheld by networks of responsible epistemic practice. Yet the risks are real: advocacy plays into a politics of knowledge where ‘advocacy research’ is a derogatory label for research undertaken to serve ‘special’, contrasted with universal (read: ‘dominant’), interests; and credibility is unevenly conferred or withheld in patterns of epistemic injustice consequent upon discounting the experiences of some would-be knowers, granting or refusing to grant their testimony a hearing, facilitating or blocking their participation in communal epistemic practices. Claiming that advocacy can make knowledge possible contests a false autonomy/advocacy dichotomy for which self-reliant ‘individual’ knowledge yields the only possible certainty, while advocacy inevitably compromises that possibility and destabilizes the putative knowledge. Like any social-epistemological project, advocacy requires ongoing vigilance, negotiation, openness to the need for corroboration, contestation, revision. It maintains a certain fallibilism in epistemic practice, yet it need not pervert the fundamental commitments of responsible knowledge-seeking. Advocacy is a vexed and vexing process for epistemologies of

mastery with their adherence to the individualism of autonomous man and its attendant image of monologic knowledge claims, uttered into a void. With social epistemology's repudiation of individualism in favour of collaborative, interactive, deliberative knowing, it is less fraught: it manifests openness to the 'counter-sentence', to negotiative processes of listening and responding; to empirical-experiential testing and retesting for confirmation or falsification. There are multiple modalities of advocacy, but charges levelled against it tend, misguidedly, to generalize from places where its imagined proximity to the worst rhetorical excesses expands to encompass all such practices.

Advocacy names practices such as those evoked in Marilyn Frye's phrase 'hearing each other into speech' (Frye 1996: 38). In such hearing—which need not be unidirectional—these practices can inform collective action intended to further certain causes, contest patterns of silencing, intervene in intransigent social orders on behalf of or in concert with the disenfranchised, lend someone's—or some group's—expertise and sometime privilege to others less well positioned to know. None of these thoughts will undo advocacy's negative implications or its dangers; but politically, these risks may have to be faced. Thus Karen Messing

47

advocates for workers so disenfranchised by material-institutional structures that they need help in speaking 'for themselves': advocacy whose credibility derives in part from her credentials and social-professional positioning, yet is true to the specificities of their situations, and available for acknowledgement or challenge from those prepared to hear and qualified to advise. Carson advocates to President Kennedy's Science Advisory Committee (PSAC) for knowledge about pesticides she has gained in field and laboratory research, and from lay testimony.<sup>27</sup> Hers was a long, bitter negotiation with and against detractors whose vested interests were at risk, and the negotiations continue after her death;<sup>28</sup> but such will be true of many advocacy-driven and other such negotiations. In short, *contra* Welbourne and Moran, giving and receiving testimony, privately (in one-on-one exchanges) or publicly, is complicated by differences in privilege, credibility, and understanding. Giving and receiving testimony is rarely a discrete event: it requires negotiating differences in privilege and complexities, in the participants' ability to communicate across differences of power, privilege, and articulateness. Messing's enquiries, and those adduced by the IDRC are cases in point.

In its responsible modalities, advocacy strives to *get at* truths operating below the surface of the presumed self-transparency of evidence. Al Gore's 'inconvenient truth' is an apt, if controversial, example.<sup>29</sup> His is unashamedly an advocacy project, urging people to face truths that are 'inconvenient' because of the losses actively acknowledging them entails. They demand radical life changes in the affluent Western/northern world with their taken-for-granted materialities and the sense of entitlement that sustains them. It is easier not to know. Precisely because of their social-political 'inconvenience', such truths require sustained advocacy if they are effectively to infiltrate the social imaginary where rhetorically fuelled counter-advocacy is actively devising strategies to gainsay them. Like

Carson, Gore has not had the last word: the controversy goes on. But his critics speak and argue in a space he has made available. At their best, advocacy practices claim discursive space for ‘subjugated knowledges’, putting them into circulation to negotiate acknowledgement; working to promote informed, emancipatory moral-political-ecological effects.<sup>30</sup> At their worst, they can do untold damage.

In these examples, ignorance maintained and sustained is a powerful force of resistance to testimonial knowledge ‘going through’. In each, there is much at stake for those required to relinquish the ignorance on which their intransigence, variously, relies. Yet as participants in the newly named ‘agnotology’ discussions show,<sup>31</sup> ignorance is more than a simple lack of knowledge: it is a multifaceted social-political phenomenon produced or sustained in ways that parallel but

48

do not replicate the production of knowledge: it is negotiated and held in place by advocacy. Its power may be consequent upon a certain stasis: a resistance within a community or group or a firmly instituted imaginary to entertain new epistemic possibilities; a vested interest in maintaining a condition of not-knowing or pervasive uncertainty about a particular subject matter; a paternalistic protectiveness against certain truths entering public knowledge or the knowledge of particular segments of the population. Ignorance thus sustained can perpetuate epistemic injustice, both testimonial and hermeneutic, in the interests of sustaining a hierarchical social order. The politics of unknowing manifests in the making and unmaking of ignorance.

In Patricia Williams's negotiations with the law review, principles of neutrality entrenched within the going epistemic imaginary blocked possibilities of knowing the salient specificities—the racial dimensions—of her situation. Moreover, presumed neutrality protected the racially unmarked from confronting their ignorance about living the differences racial difference makes, or acknowledging it in any but a blandly liberal register for which ‘we’, whoever we are, are all alike beneath such allegedly insignificant differences as skin colour, gender, race, ethnicity, disability, and other such hermeneutical and testimonial injustice-fuelling specificities.<sup>32</sup> Given the constitutive part assumptions of universal human sameness play in the received epistemologies and moral-political theories of Anglo-American philosophy, it is apparent that change can be effected and multiple modalities of ignorance dislodged only at the level of the instituted social imaginary. These are no individual, punctiform errors in perception, interpretation, or judgement to be corrected with a simple dose of counter-evidence.

Forces of carefully preserved ignorance are at work in the epistemic climate where Carson, Messing, and the IDRC workers sought acknowledgement for their findings. An advisor warned Carson: ‘Facts will not stand in the way of some confirmed pest control workers and those who are receiving substantial subsidies from pesticides manufacturers’;<sup>33</sup> Messing faced entrenched stereotypes of women's ‘nature’, invoked to condone a systemic ignorance of the specificities of their situations which would have required action, at some cost, by employers; the IDRC faced aid-distribution policies

based in sustained ignorance of local epidemiological facts, derived from policies not tested in their ecosystem specificity. Curiously, in the light of resistance to advocacy, received theory and practice are sustained and safe-guarded by persistent advocacy *against* according acknowledgement to the local testimony that ultimately made knowledgeable responses possible in each of these examples. I mention these issues not to level a

49

*facile tu quoque* counter-argument against advocacy detractors, but to reinforce my contention that, in social epistemology the autonomy/advocacy dichotomy is less readily tenable than it has been in the rhetorical spaces where ‘advocacy research’ is denigrated for its alleged commitment to serving ‘special’ interests, and is contrasted with impersonal, universal truth-seeking enquiry which, too often, silently beneath its claims to neutrality winds up serving dominant interests, after all. The epistemic injustices it perpetuates would be impossible to catalogue.

50

### 3 Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time

Miranda Fricker

#### 1. DIMENSIONS OF SOCIALIZATION

We tend to think of the socialness of social epistemology largely in terms of a lateral expansion across social space. The expansion shifts the philosophical focus from the lone individual of so much traditional epistemology—the individual who wonders whether he really knows this is his hand before him, and so on—to his relations with his fellow subjects, his epistemic interactions with them, even his epistemic interdependence with them. The interest in epistemic interdependence brings divisions of epistemic labour centre-stage (as the explosion in the literature on testimony in recent years bears witness) and further establishes a recognized theoretical space for insights about how justification (for instance, justification for a scientific theory) might be dispersed across a whole epistemic community, with the consequence that it makes sense sometimes to regard that whole community as the subject of the knowledge, and perhaps no individual at all.<sup>1</sup> This kind of socialization of epistemology, then, brings with it a new, less individualistic conception of epistemic subjects. No longer conceived as lone individuals whose interactions with other individuals are epistemically incidental, we think of them as fundamentally, naturally, placed in relations of epistemic interdependence. Let us call this socialized conception of epistemic subjects, the *Abstracted Social Conception*. It marks the anti-individualist moment in epistemology.

The conception is so labelled because it remains highly abstracted—appropriately for certain purposes. The social relations in which epistemic subjects are conceived as standing are relations between finite knowers and enquirers conceived as bearers of reasons, producers of evidence, seekers of information, conveyors of knowledge, and so on. But these knowers and enquirers are

signally not conceived as standing in relations of social identity and power—the conception abstracts away from that grade of social detail. Imagine a painting that represents a plurality of figures engaged in some form of interaction; but where the figures are painted abstractly, leaving it indeterminate whether they are young, old, female, male, black, white, rich, poor, dominant or subordinate, and so on. The Abstracted Social Conception is like that, and so it does not raise the question (being unequipped to answer it) of how these social details might be affecting the interactions we see represented, or whether there are other kinds of interaction going on that relate specifically to identity and power. In order to be open to such questions, we need a conception that is more like a painting that represents figures in their full social colour; we need to make available a *Situated Social Conception* of the epistemic subject. It

should not replace but rather complement the more abstracted conception, for both conceptions have a role. Categories of identity and power are only relevant for certain sorts of philosophical question, after all, and those operating with the Abstracted Social Conception are not on the whole aiming to raise them. (On a historical note, however, it must be said that the Abstracted Social Conception doesn't easily *allow* such questions to come into view as genuinely epistemological questions, so that questions involving identity and power have tended to appear as questions for the sociology of knowledge alone, which is at least part of the story why traditional epistemology remained so thoroughly asocial for so long.)

The two conceptions are not really all that far apart, however. A suggestive distinction often featured in the literature is that between layman and expert. One of the things that makes this distinction interesting is that it can be taken as a purely epistemic distinction, so that we conduct our debates about it in somewhat rationally idealized terms, yet it can also be taken as a distinction between two social identities where there are relations of power that hold between the two parties. A layperson's relationship to an expert does not of course have to involve any significant power relation, in the sense of a relationship of power that affects their interaction or impinges on the rationality of their exchange. As John Hardwig has pointed out, in very large and complex scientific projects where there is a marked division of epistemic labour among the contributing scientific communities, it makes sense to regard all the contributing individuals as at once experts vis-à-vis their own contribution and laypersons vis-à-vis the contributions of others.<sup>2</sup> In such a scenario, what relations of identity and power there may be between different groups of scientists may not give rise to any epistemologically compelling issues. But, then again, they might. In a scenario where one set of contributors happens to enjoy more professional esteem than others (perhaps, simply, they are operating under the auspices of an especially powerful institution) there could easily arise the sort of mingling of power with norms of enquiry that is not easy to disentangle. When this

happens, it becomes part and parcel of the scientific enquirer's requisite epistemic virtues—possessed either by individuals or possibly only by the community—to reliably succeed in spotting research decisions that are too much driven by professional or institutional power. When social epistemologists talk of experts and laypersons, then, they are already flirting with a more fully socialized conception of epistemic subjects than the Abstracted Social Conception itself allows for.

If we want epistemology to account for the *human* epistemic predicament, then we need to have available a conception of epistemic subjects as required to overcome or negotiate certain entanglements of reason and power, because it is an essential feature of human enquirers that they operate in a context in which such entanglements can arise. This is the source of the value of the Situated Social Conception, which conceives epistemic subjects and their interactions as situated in a context of social identity and power, and so makes

visible the influence of these factors on our epistemic interactions. For any given project in social epistemology, then, we need to be reflective about which conception suits our philosophical purposes—reflective, that is, about which degree of abstraction is appropriate for the issues we want to bring out. Simply sticking to the standard Abstracted Social Conception may occlude ethical and political aspects of epistemic practice that are worth our attention; then again, sustaining a Situated Social Conception will be pointless if relations of identity and power are irrelevant to the issue we are pursuing. It's a judgement about horses for courses, so, as a matter of good philosophical method, we need to have the different options reflectively available to the philosophical imagination. The picture of epistemic subjects presented by the Situated Social Conception is less abstract than that presented in the Abstracted Social Conception, but it is still an abstraction, as befits the philosophical purpose. It represents epistemic subjects not in their personal detail—obviously not—but as variously instantiating one or another more or less complex social type. If the Abstracted Social Conception marks the moment of rebellion against excessive individualism in epistemology, the availability of the Situated Social Conception marks our graduation beyond the naivety of compulsory rational idealization.

I have argued elsewhere for the importance of the Situated Social Conception for certain philosophical purposes, and in particular I have argued that there are issues of justice and injustice in our everyday epistemic interactions—our testimonial interactions and our practices of social interpretation—which cannot come to light unless we adopt that more fully socialized conception.<sup>3</sup> For present purposes, however, the Abstracted Social Conception is appropriate, as my aim is to show how expanding not only across social space but also across *time* can be a powerful epistemological resource. We should distinguish between two sorts of temporal expansion: expansion across real time (including historical time), and

expansion across the quasi-fictional time that is at work in genealogical method. I shall make a case for the philosophical fruitfulness of expanding over genealogical time, and my specific aim will be to show how the genealogical method can support and augment certain socializing arguments against scepticism. The genealogical story I shall use is that given by Edward Craig (1990) in his book, *Knowledge and The State of Nature* (hence the appropriateness here of the Abstracted Social Conception, for in so far as there are any social types in the State of Nature, their social identities do not figure in the explanatory purpose that this particular genealogy aims to achieve).

I shall make my case by reference to Michael Williams's diagnostic engagement with scepticism, in which he crucially employs a Default and Challenge model of justification. And I will develop three key aspects of Craig's 'practical explication' of the concept of knowledge so that they may be seen to resonate positively with Williams's epistemological picture: the admixture of internalist and externalist features (3.1); the



proto-contextualism (3.2); and, finally, the distinctively genealogical anti-sceptical impetus (3.3). In this way I aim to support and augment the socialized anti-sceptical case mounted by Williams. I also aim thereby to illustrate the philosophical productiveness of expanding epistemology not only laterally across the social space of other epistemic subjects, but at the same time vertically in the temporal dimension.

## 2. DEFAULT AND CHALLENGE: SOCIALIZING JUSTIFICATION

In 'Responsibility and Reliability'<sup>4</sup> Michael Williams (2008) explores the anti-sceptical impetus of a certain model of justification that is found in Robert Brandom's work under the name of Default and Challenge. It is an entitlement conception of justification, according to which we may assume our faculties are functioning correctly so long as there are no reasons to suspect otherwise. Williams's paper has two aims. Firstly, he aims to incorporate reliabilist insights within a fundamentally deontological framework, where the key reliabilist insight he has in mind is that many accounts of knowledge incorporate exaggerated, over-intellectualized conceptions of what it takes to be epistemically responsible, which simply do not square with the spontaneous and unreflective character of our most basic forms of knowledge, notably, perception. If there is to be a satisfactory responsibilism that presents a unified account of knowledge, it will have to avoid such intellectualism. And, secondly, he wants to show how such

54

a responsibilism can, by the same token, deflect scepticism. The problems of intellectualism and sceptical challenge, he argues, have a common solution, for one and the same excessively internalist, mentalistic conception of justification is their common root.

Williams traces a dominant internalist conception of justification back to Chisholm, and the model he finds in Chisholm's writing is one in which justification constitutes a kind of 'positive authorization' which, in Chisholm, is linked to a foundationalist structure with error-proof sensory experiences at the bottom, so that the whole structure is designed in the foundationalist style to stave off sceptical challenge. Looking to Sellars's critical analysis of this sort of view and its dependence on notions of the Given, Williams argues that the positive authorization conception seriously exaggerates what is needed to vindicate the idea that epistemic subjects achieve justification by acting *in the light of* epistemic rules, as opposed to merely conforming to them. We can achieve a picture of subjects acting in the light of normative rules, without being compelled to add that rules should be construed as imperatival in form, or that justification flows upwards in the system from a foundation of error-proof self-addressed 'reports' of experience. Williams continues in the Sellarsian idiom by taking up a distinction Sellars makes between 'ought-to-do' rules, which are imperatival in form, and 'ought-to-be' rules, which are not. These so called 'ought-to-be' rules effectively set conditions of entitlement in the Default and Challenge mould. In Williams's example:

For me to see, and not merely think that I see, that there is a rabbit in the garden, all sorts of conditions *must be* met. Some concern me: I must be of sound mind, paying attention, capable of recognizing what is going on, and so forth. Others concern the object and its situation: the animal has to be a rabbit and not a stuffed toy, the light must be good enough to make out the shape of the dark patch in the middle of the lawn, and so on. If these conditions are not met, I won't be in a position to see that there is a rabbit in the garden. (Williams 2008: 12–13)

What Default and Challenge achieves for us is the desired admixture of internalist and externalist insights. In order to count as acting in the light of a rule (in order to count as epistemically justified) the well-trained subject might, depending on the context, need only to be counterfactually sensitive to lapses in the conditions required for taking the deliverances of her faculties for granted. As Williams puts it, 'Our acceptance of an ought-to-be "rule" consists principally in our disposition to acknowledge the exceptions, and to respond appropriately' (2008: 21). In sum, the well-trained subject may take her sensory experiences at face value, so long as there are no reasons not to. In doing this she is following rules of justification, acting in the light of them but not self-consciously. Thus the externalist aspect of Default and Challenge that sets it apart from any 'positive authorization' model. Yet if appropriately challenged, she does have a standing obligation to produce a justification, and if she cannot, then she is revealed as

55

lacking entitlement to her belief.<sup>5</sup> Thus the internalist aspect of Default and Challenge that qualifies Williams's position as a form of responsibilism.

Williams also argues that seeing justification as conforming to a Default and Challenge structure can help fend off scepticism. He invokes the distinction between Agrippan and Cartesian forms of sceptical challenge. Agrippan scepticism imposes an endless demand for further justifications, so that it threatens to expose either regress or circularity in the series of justifications we may offer, or else a plainly unjustified assumption somewhere in our reasoning. Cartesian scepticism is characterized as exploiting issues of underdetermination by positing sceptical scenarios which he claims, for all we know, we might be in.<sup>6</sup> Williams's focus, however, is on the Agrippan style sceptic, and he argues that the Agrippan is implicitly committed to the familiar, mentalistic and so excessively internalist model of justification that conceives being guided by norms or rules as always a matter of self-conscious obedience to self-addressed imperatives—positive authorization. What the Default and Challenge model furnishes is an account of justification—entitlement—that makes no such requirement. By contrast, what Default and Challenge obliges the individual subject to do is something negative: don't take your experiences at face value if the default condition is lapsed. A subject who is entirely successful with respect to that negative task may well not be able to answer the Agrippan sceptic—but so much the worse for the sceptic. The standing obligation to come up with reasons when challenged only holds for challenges to which one's interlocutor is entitled. Brandom, introducing the label 'Default and Challenge', puts the point like this:

Claims such as 'There have been black dogs' and 'I have ten fingers' are ones to which interlocutors are treated as *prima facie* entitled. They are not immune to doubt in the form of questions about entitlement, but such questions themselves stand in need of some sort of warrant or justification. Entitlement is, to begin with, a social status that a performance or commitment has within a community. . . The model presented here has what might be called a *default and challenge structure* of entitlement.

(Brandom 1994: 177)

The Agrippan sceptic, then, is presenting inappropriate challenges, and so our failure to meet those challenges signifies nothing. The sceptic is thus revealed as missing the point, for she tries to compel the individual subject to dig deeper and deeper into his individual epistemic resources, furnishing reason upon reason for his belief—but this is simply the wrong place to look for his status as justified in

56

believing what he believes. The Agrippan demands to be shown a justificational stopping point somewhere in the depths of the individual subject, and her mistake is that justification is not to be found deep in the individual but rather on the surface of something irreducibly social, namely, the subject's ability to meet the challenges properly brought in that context by others in the epistemic community.

Williams focuses on the mentalistic nature and extreme internalism of the model of justification that both polarizes reliabilism and responsibilism and hands the sceptic a stick to beat us with. But I think we should most of all emphasize its *individualism*, for it is the individualism that underpins both the mentalistic and the excessively internalist character of the mistaken model of justification that Williams rightly diagnoses in the sceptic. Given a general assumption of epistemic individualism, one easily sees how it can seem natural to assume that the individual is the source of all justification for her belief, and that the place to look is (where else?) in her mental states. Thus the mentalism. Further, if even the subject herself cannot find a justification in her psychology, then how can she count as possessing a justification at all? Thus the extreme internalism. I think, then, that once we focus on the anti-sceptical energy that Default and Challenge clearly contains, we find that energy to be derived most fundamentally from its *sociality* (and so, to relate back to the terminology I introduced earlier, from its implicit insistence on the Abstracted Social Conception of epistemic subjects). For it is the sociality that effects the crucial shift in rational obligation away from the lone individual and into the epistemic social body as a member of which the individual subject must function. It is of course individuals who bear the responsibility of justification for their beliefs, but the point is they can only live up to this responsibility in virtue of their participation in a social practice of proper challenge. (I alone bear the responsibility for paying my bills; but the question of what I owe is settled as a matter of social practice.) It is this social dispersal of justificatory labour that relieves the individual believer of the burden of accessing the kind of justification that the Agrippan sceptic demands.

The Agrippan, who presses and presses for ever more justifications, is thus revealed as making a profound mistake at both the level of epistemic practice and the level of epistemology. She fails to adhere to socially established norms of challenge, a mistaken practice that exposes her false theory of justification. But what of the Cartesian sceptic? Williams thinks the Cartesian, whose signature is of course the sort of madcap sceptical scenarios we all know and love, cannot be confined in the same way, but only indirectly 'by way of showing how we can legitimately set sceptical problems aside' (Williams 2008: 26). Why is this? If the context sensitive norms of challenge can reveal Agrippan challenges as mistaken, then why not Cartesian challenges too? Now that we have identified the sociality of Default and Challenge as fuelling the anti-sceptical work, is it not evident that the Cartesian invocation of sceptical scenarios is a style of challenge every bit as inappropriate as the Agrippan? The Cartesian, too, seems to press

57

his case to the individual knowledge-claimant in a way that is not sanctioned by socially established norms of Default and Challenge—the demand to rule out that one is a brain in a vat, or whatever it may be, is obviously outrageous. I think this does have some bite as an argument against the Cartesian. He, no more than the Agrippan, can just assume that his challenges are appropriate—on the face of it they certainly are not.

It remains true, however, that the Cartesian has a card up his sleeve that the Agrippan lacks, in as much as the Cartesian's challenge plays specifically on the idea of underdetermination in order to spook us with the possibility that our beliefs are not justified at all, on the grounds that they are no more justified than the alternative 'theories' of the world cooked up in the sceptical scenarios. This means the Cartesian can be construed as making a challenge not necessarily to the individual knower in any way that violates the insights of Default and Challenge, but rather as making a higher level challenge in respect of what we think we achieve by adhering to our precious norms of appropriate challenge—not knowledge, he exclaims, not even justified beliefs! Admittedly, then, revealing the Cartesian as having an excessively individualistic theory of justification that should be replaced by Default and Challenge still leaves the spooky radical sceptical possibility rather as it is, which is why other arguments (against 'epistemological realism' and for an antidote contextualism) need to be brought in. Nothing in Default and Challenge, after all, directly addresses issues of underdetermination; whereas, by contrast, the favourable identification of Default and Challenge as a rival account of justification reveals the Agrippan as entirely driven by the mistaken assumption that justification requires an individual capacity to respond to an indefinite series of challenges, without lapse into regress, assumption, or circularity. Even while I want to insist that the socializing move inherent in Default and Challenge does win a trick against the Cartesian sceptic, then, still it is clear that it cannot on its own beat the Cartesian's hand in the way it beats the Agrippan's.

Accordingly, in respect of the Cartesian challenge, Williams looks to the central anti-sceptical argument of *Unnatural Doubts* (1991), namely, the argument that the Cartesian

sceptic is committed to epistemological realism and epistemological realism is false. The Cartesian sceptic's challenges concern something called 'knowledge of the external world' or 'empirical knowledge', as if these were respectable theoretical categories; but they are not. They are far too internally diverse to be so regarded, and in fact have no more integrity than a category such as 'knowledge of things done on a Wednesday'. Crucially, they are too internally diverse in terms of the kind of justification that is required—something we may express in terms of the Default and Challenge model by saying that the norms of appropriate challenge vary from context to context. The Cartesian sceptic may possibly go in for his peculiar style of challenge in the strictly 'epistemological context', but not in other contexts. To do so would, as ever, constitute a mistake at the level of norms of Default and Challenge governing our epistemic practice,

58

but more importantly perhaps, it would be an enactment of the false piece of theory that is epistemological realism. The Cartesian sceptic wants to move from (i) discovering that, in context C, knowledge is impossible, to (ii) discovering (in context C) that knowledge is impossible<sup>7</sup>, but if he can only make that move by way of the false doctrine of epistemological realism—a staging post which would effectively privilege the so-called 'epistemological context' over all others—then the move is blocked and, *qua* sceptic, he is confined to the study. That is, his eccentric style of justificational challenge is confined to the context of enquiry that is peculiar to a certain style of epistemology.

Thus Williams's fascinating anti-sceptical case. I have so far discussed (and slightly elaborated) the use he makes of Default and Challenge against scepticism in 'Responsibility and Reliability'; and I have recalled (as he does) the contextualist position he first argued for in *Unnatural Doubts*, where it functions as the antidote to the sceptic's epistemological realism. My chief purpose here, however, is to make a case for the expansion of our philosophical conception of epistemic subjects and their activities along a certain temporal dimension, namely, the genealogical temporal dimension. So how might a genealogy of knowledge help bolster and augment Williams's anti-sceptical case?

### 3. EXPANDING ALONG THE TEMPORAL DIMENSION—GENEALOGICAL TIME

In his *Knowledge and The State of Nature* Edward Craig gives what I'm calling a genealogical account of knowledge. That is, he tells a State of Nature story about why we have the concept of knowledge—a 'practical explication' of that concept. He envisages a minimally social epistemic community—an abstraction of any real human community, though one that non-accidentally resembles what a real early human community might have been like in respect of its social simplicity and its hand-to-mouth relation to basic human needs and dangers. The basic epistemic needs that define the State of Nature are, first, the need for enough truths (and not too many falsehoods) for other sorts of basic needs—principally survival needs—to be met. A community that survives in the State of Nature must operate with sufficient truths to hunt and/or forage for food, take care of the young, avoid predators, deal with the dead, and so on. That first epistemic need

immediately gives rise to a second: the need to realize the epistemic and practical advantages of pooling information. Why rely only on one's own eyes and ears when you can benefit from the eyes and ears of others? From where you're standing you may not be able to see if the predator is coming, but that person up in the tree might, and this exemplifies the

59

fundamental practical pressure to stand in cooperative epistemic relations with fellow enquirers. Finally, this second epistemic need spontaneously gives rise to a third: the need to distinguish good from bad informants, so that it is indeed information that gets shared and not misinformation or disinformation. Human beings, however described, are fallible—hence the risk of misinformation. And human beings in the State of Nature, as anywhere else, operate under pressures (such as competition for resources) that create motivations for deception—hence the risk of disinformation. Distinguishing good informants is indeed an essential capacity.

This trio of fundamental epistemic needs generates a certain point of view for our social epistemological project: the point of view of the enquirer. This is notably different from the point of view normally taken up in epistemology, namely that of the examiner; a point of view typified by the epistemologist's remove from the actual business of enquiry in order to debate about whether some candidate knower really qualifies.<sup>8</sup> The particular need to distinguish a good informant as to a given question whether  $p$  is a need had only by someone who doesn't know whether  $p$  but wants to. Accordingly, as we construct the epistemic State of Nature, we find that ignorance and the desire to make it good with good information emerge as our basic epistemic state. In this sense, 'Who knows whether  $p$ ?' is our most basic epistemological question, a question that presupposes the possibility of knowledge. How far this broad anti-sceptical presupposition has any argumentative force depends upon how convincing the overall story is in terms of its explanatory power. In so far as the State of Nature construction provides a convincing explanatory story about why we have, of necessity, the concept of knowledge, then so far may it turn out to give genuine independent support to the idea that sceptical questions are parasitic on there being a functional epistemic practice in which knowledge is possessed and, in particular, shared or 'commoned'<sup>9</sup> in an epistemic community. (I shall return to this in section 3.3.)

So how does Craig's genealogy explain the advent of the concept of knowledge? We have seen that the enquirer needs to distinguish good informants. A good informant is someone who: (1) is likely enough in the context to be right about whether  $p$ , (2) is communicatively available and open (including sincere), and (3) bears indicator properties so that you can reliably recognize that (1) and (2) are satisfied.<sup>10</sup> In Craig's story, indicator properties will be a mixed bag, but might standardly include properties such as having been looking in the right direction at the time, or having a good track record.<sup>11</sup> Craig's thesis is that the

60

constructed concept of the good informant constitutes the core of our actual concept of a knower. As we might put it, the status of being a knower starts life as the status of being a good informant. The two concepts are not coextensive of course: there can be knowers who are not good informants, for instance because they lack the requisite indicator properties, or because the indicator properties (being only reliable) unluckily mislead on that occasion. But Craig's proposal is that the functional origin of the concept of knowledge is to identify good informants; and that thereafter the constraints of recognizability and communicative openness gradually become relaxed, as certain more sophisticated uses we come to make of the concept (such as referring to knowers we cannot ourselves recognize) pressurize it in the direction of 'objectification'—that is, of referring to something that exists independently of our powers of recognition. Such pressures explain how we come to think of knowledge as something another person can possess, even if we can't recognize it, or they aren't coming out with it.<sup>12</sup>

### **3.1. The Original Synthesis of Internalist and Externalist Insights**

According to Craig's genealogy, then, we start to operate with the concept of knowledge, of necessity, because at the core of that concept is something that meets the absolutely basic epistemic need to pick out good informants. Now how does all this help the socializing anti-sceptical case that Williams builds on Default and Challenge? One of the key anti-sceptical features of Default and Challenge is that it achieves a desirable combination of internalist and externalist features. I think we can see how this is explained and so reinforced if we look closely enough at epistemic practices in the State of Nature. In the first instance, the practice of pooling information in the State of Nature features people spotting others as good informants before asking them for information. But we can see how the basic need for good information also drives a modification of that practice; namely, asking candidate good informants the question to which we want the answer, and then, once they have responded, quizzing them as to their reasons. The capacity to give reasons for what one asserts is a supremely important indicator property, not discussed by Craig. The person who asserts but does not know may be suspiciously fuzzy on her reasons; the person who asserts what he knows to be false may be suspiciously

unconvincing when he presents purported reasons for his pretend belief. The ability to satisfy this kind of challenge is a key indicator property of good informants.

Being able to supply a justification when challenged is not, however, a necessary condition of being a good informant. Rightly not, for what primarily matters to the enquirer is simply that the good informant comes out with the truth on demand; not that he comes out with his reasons as well. Given that the enquirer can spot a good informant to her own satisfaction, she will just take the information and not bother to quiz him further about his reasons. However, this basic practice established, we can immediately see how quickly an informant's capacity to give reasons assumes importance, for it is highly desirable in a good informant that he be able to produce reasons when challenged,

owing to the fact that this may be by far the best indicator property available to the enquirer. The same point applies individualistically too, for enquirers in the State of Nature will often be relying on the deliverances of their own faculties, and are best construed as entitled to trust them unless they have some reason not to—a foggy day; a foggy memory.<sup>13</sup> (Here we glimpse the nascent Default and Challenge structure of justification emerging.) On these occasions, a certain challenge to self is in order, which amounts to a demand for an after-the-fact indicator property that one is likely enough in the context to be right about  $p$ . The ability to produce a satisfying reason is the prime case of such an indicator property.

The importance of this capacity to come out with reasons when appropriately challenged, combined with the fact that it is not one of the conditions of qualifying as a good informant, explains what underpins the desired admixture of internalist and externalist features that Williams aims to achieve. The picture in the State of Nature is fundamentally externalist—what matters is simply that good informants come out with the truth—but then we quickly come to see the origin of internalist intuitions about knowledge. On the story I am urging here, we agree with Craig that the good informant's capacity to access his reasons is not at the core of the concept of knowledge. But, we add, nor is it a merely peripheral feature, resulting from some historical accident in how we use the concept. Rather, the capacity to access one's reasons features in a layer of content that is close to the core of the concept of knowledge. The importance of the capacity to produce reasons in support of what one believes flows immediately from the basic method of identifying good informants that does constitute the core. I think this is a good way of substantiating the two-sided thought that it is close to conceptually impossible that a human being who lacked the general capacity to come out with reasons for her beliefs could count as a knower (or even a believer); but being able to come out with one's reasons is not thereby required in every instance of knowledge, and so not a necessary condition.

This combining of externalist and internalist features of the practice of justification on the part of good informants in the State of Nature echoes and substantiates the internalist-externalist combination we find in Default and Challenge. On that model of justification, the subject can have knowledge even if she has taken the deliverances of her faculties entirely at face value and cannot produce any positive reason for her belief. If faced with an eccentric challenge, she may only be able to assert that it never occurred to her to wonder; she may even be a bit thrown by the fact of the challenge and by her own bewildered reaction. And yet, according to Default and Challenge, the fact that the default did indeed hold is sufficient. The kind of responsibilism Williams's arguments are designed to achieve, I think, is one that allows externalism vis-à-vis the question whether the default of entitlement holds, so that there is no blanket requirement that the subject be reflectively aware that it holds; yet a modest internalist constraint when it comes to the subject's obligation to respond to contextually appropriate challenges. My suggestion has been that a responsibilism of that combinatory sort finds explanatory support in the



genealogical approach, owing to the non-core yet close-to-core role that an ability to produce reasons plays in the State of Nature.

### 3.2. Practical Origins of Contextualism

What about contextualism?—a basic version of which might be considered part and parcel of the Default and Challenge model of justification.<sup>14</sup> I suggest that this too finds an origin in the State of Nature. We have seen that the enquirer is looking for someone who is crucially *likely enough in the context to be right as to whether  $p$* . This presents an explicitly contextualist picture of its own, according to which what it takes to be a good informant—what it takes to play the social role at the core of knowing—alters from context to context. For instance, if the stakes are very high, the good informant will be almost certainly right about  $p$ ; if they are not so high, and/or if one needs to act sooner rather than later, she might count as likely enough to be right about  $p$  just by having a half-decent track record and being the only candidate good informant in the vicinity. The State of Nature, then, explicitly imposes the aspect of contextualism that Williams has elsewhere labelled ‘economic’:

If it is important to reach some decision, and if the costs of error are fairly low, or if we gain a lot by being right and lose little by being wrong, it is reasonable to take a relaxed attitude to justificational standards. If the costs of error are high, more demanding standards may be in order. The opportunity costs of further inquiry can also be relevant.

(Williams 2001: 161)

The proto-contextualism contained in the State of Nature not only signifies a constraint on who can be a good informant, it imposes a constraint on the enquirer too. If, in our imagined scenario in which the enquirer has already asked a candidate good informant for information, the informant has told her something, and she is challenging him for his reasons, then her challenges must be appropriately geared to the context. She is looking for an after-the-fact-of-utterance indicator property, and this means that if she were to press him for a justification that exceeded the contextually required level of probability, this would mark a dysfunction in her epistemic conduct from her own point of view as an enquirer. Basic practical concerns generate the context sensitive norms of Default and Challenge in the State of Nature, so that enquirers who demand reasons above and beyond those appropriate to the context are making a mistake at the level of the (emergent) norms of proper challenge. All in all, the State of Nature presents us with a range of practical contexts such that in some contexts the enquirer will demand very little by way of an informant's likeliness of being right, whereas in other contexts she will demand a good deal more, according to her practical needs and interests—according, that is, to what it takes for the informant to be likely enough in the context to be right about  $p$ .

This aspect of context is the only one we find explicitly stated in the formulation of our State of Nature scenario, as a condition of being a good informant, though the other determinants of epistemic context can be built in if one is so inclined. Besides the

economic, Williams cites four other aspects of context: (1) the intelligibility of error—we cannot make sense of the possibility of error except against a backdrop of getting it right; (2) issues of methodological necessity—in order to pursue any question we must take some others for granted; (3) ‘dialectical’ issues—default entitlements are lost and gained according to the ongoing movement of evidence; and (4) ‘situational’ issues—claiming knowledge commits us to the objective well-groundedness of our beliefs (either a default entitlement holds or it doesn't). Taking as an example the issue of methodological necessity, people in the State of Nature definitively don't do history, so they don't have to set aside questions about the reality of time in order to do it; but they may well have to establish whether a certain water source is clean or not, and so they do have to set aside questions about the reality of the external world.

I believe something similar can easily be said in relation to the other aspects of context that Williams argues for, but I do not want to set much store by it. I dare say an advocate of another style of contextualism could construct a similar support in the State of Nature for his particular brand of the view too. Only the contextualism we find explicit in the formulation of the State of Nature scenario should define the proto-contextualism we regard as proper to it, and that is the condition of being a good informant that requires her to be likely enough in the context to be right about *p*. My point is that the practical pressures that generate this proto-contextualism lend genealogical explanatory support to the

contextualism that Williams embraces, even while we can see that they might also be able to support distinct forms of contextualism in which the ‘economic’ element was differently embedded. That is entirely as it should be, for the State of Nature scenario purports to contain only the necessary features of our practices. We should not try to find anything besides the core elements of our actual justificatory practices represented in the State of Nature scenario, and so should not hope to extrapolate anything more than a generic proto-contextualist commitment.

### **3.3. The Un-Originality of Sceptical Challenge**

In this exploration of the practical pressures that generate contextualism in the State of Nature, we already begin to see how the genealogy of knowledge provides independent support for the anti-sceptical purpose to which Williams puts his own contextualist position. There are no sceptics in the State of Nature—survival requires taking some people as knowing things one needs to know, and that entails accepting the possibility of knowledge. This underpins my earlier suggestion that Williams's charge against the Agrippan sceptic—that he behaves in a contextually inappropriate manner—does have some force of its own against the Cartesian sceptic. In the State of Nature, it doesn't matter in which style sceptical challenges are made. The fact that any such challenges make demands that exceed what it takes for the informant to be likely enough in the context to be right about *p* means the sceptic will fail to identify good informants that are staring him in the face, and will lose out on knowledge as a result. This means that not

only the Agrippan demand for ultimate justification but also the Cartesian demand that we meet the challenge of underdetermination can only be a mistake.<sup>15</sup> The State of Nature, then, explains the commonsense idea that no one can *basically* be a sceptic. They must be enquirer first, and sceptic second; someone committed to the practical possibility of knowledge first, and committed to undermining that possibility second. This of course leaves room for the idea that there may yet be a context in which it *is* appropriate to mount sceptical challenges. The present argument shows only that there is no such context in the State of Nature.

This accommodating thought prompts exposure of the other respect in which we can see the genealogy of knowledge lending independent support to Williams's anti-sceptical strategy. The point of extending our philosophical conception over the semi-fictional time in which genealogical narratives are set, is

65

that it provides an invaluable way of *relating* core features of a concept to non-core and peripheral features. I know of no other philosophical method that provides the opportunity to relate original, necessary features of a concept to less basic, more contingent—historically contingent—features.<sup>16</sup> Accordingly, it delivers an entirely different image of concepts than that issued by the analytical method. The analytical ambition and attendant philosophical imagination generates an image of concepts as like molecules, ready for their different elemental components to be separated out by the philosopher acting in his capacity as conceptual chemist. Genealogical method, by contrast, brings with it an image of core and periphery, or kernel and outer layers—the kernel presents necessary features of the concept, and the outer layers increasingly contingent historical features. These layers may be separated out from the kernel by the philosopher acting in his capacity as something more like conceptual historian. The necessity of the core features stands or falls with how convincingly the story passes muster as a pure construction out of nothing but absolutely basic needs. If, however unwittingly, one includes a contingent feature in the State of Nature scenario, perhaps to suit one's philosophical purpose, then, clearly, the story will lack force. No doubt every story of origins should be accompanied by something of a health warning, for it surely is all too easy to portray the State of Nature in one's philosophically preferred image. As Foucault scoffingly warns us, echoing what he takes to be Nietzsche's own warning about the philosopher's fantasy of the origin:

History is the concrete body of a development, with its moments of intensity, its lapses, its extended periods of feverish agitation, its fainting spells; and only a metaphysician would seek its soul in the distant ideality of the origin. (Foucault 1984: 80)

But Foucault is wrong if he equates all origins stories with ahistorical fantasy. One of the great virtues of the State of Nature method is precisely its separation of features of a concept that bear the necessity of the origin, from features that are more or less contingent matters of history. Far from demanding some dubiously 'metaphysical' faith in the necessity of any given practice, it gives us a concrete way of assessing the

plausibility of any such claim. When confronted with a genealogical claim that such-and-such a practice is humanly necessary or ‘original’, we must ask, Does this practice really qualify as an indispensable feature of human life? Is it really required to meet absolutely basic human needs? Can we not imagine a recognizably human society without it? Genealogical claims of necessity, then, are grounded not on dubiously ‘metaphysical’ ideas or articles of analytical faith, but rather on something fundamentally practical—the *practical human necessity* of the materials used to construct the State of Nature

66

scenario. If we can find nothing contingent about the posited original human need to pool information—if, that is, we cannot make sense of the idea that there could be a recognizably human society absent this most basic form of epistemic cooperation—then the thesis that identifying good informants comprises the kernel of the concept of knowledge possesses significant force. Foucault seems to regard it as a misuse of the genealogical method ever to use it to identify human necessities. But if it were in the very nature of genealogical method that it could only properly be used to find contingency and never necessity, then it would be a lousy, because indiscriminating, method. Foucault’s blithely expressed wisdom that when it comes to concepts, all is history and nothing origin, is mere prejudice.

Craig’s State of Nature story reveals that if there is a context of enquiry in which sceptical challenges are appropriate, still there are nonesuch in the State of Nature. And from this we have drawn the inference that sceptics must be enquirers first and sceptics second (or, *basically* enquirers, and sceptics only superficially). This, in its own right, blocks the sceptic’s colonizing move from (i) discovering that, in context C, knowledge is impossible, to (ii) discovering (in context C) that knowledge is impossible. What blocks scepticism here is the *genealogical primacy* of practical, knowledge-permitting, contexts of enquiry. These knowledge-permitting contexts are the contexts in which the core of the concept of knowledge is dramatized in practices of good informing. In this sense, knowledge-permitting contexts figure at the core of the concept of knowledge; indeed they exhaust it, for there are no other contexts in the State of Nature. Thus the possibility of knowledge is prior to the possibility of sceptical challenge in the special sense that can only be supplied by imaginatively stretching our concepts of knowledge and justification across genealogical time: even the sceptic cannot escape the cognitive functionality of the origin, for that scenario is still with us, still supplying the core of what it is for us to know.

As in Williams’s irenic anti-sceptical strategy, this may still leave some room for a confined practice of sceptical challenge—it is only in the exclusively practical contexts provided by the State of Nature that sceptical challenge is obviously never a proper challenge. In the real-time practices of Default and Challenge there may possibly remain a context in which sceptical challenge is appropriate, so that knowledge is not possible in that context. But such an ‘epistemological context’, if there be such<sup>17</sup>, exploits only aspects of the concept of knowledge that are way out on the periphery. And our

genealogy has exposed the historical contingency of merely peripheral practices, so that the sceptic may at any time find the locks to his study have been changed and that it is now being put to a different philosophical use. In this way genealogical time

67

has implications for historical time: while the diagnosis of the sceptical urge goes deep in philosophy, our genealogy of knowledge reveals that the question of the propriety of sceptical challenge does not go deep. It simply rests on how much nurturance we continue to give to the context of enquiry in which sceptical challenge is deemed appropriate. That is, it rests on something social within the philosophical community: namely, how far we continue to respond to sceptical challenge in the epistemological context as justified challenge, how long we continue to sustain something called the 'epistemological context'. Perhaps only a satisfactory theoretical diagnosis can catalyse and justify the historical shift that would render sceptical challenges inappropriate in all contexts; and that is surely something to which Williams's arguments make a significant contribution. For my part, I have tried to show that the genealogical approach contributes an independent diagnostic strategy, which can be seen to support and augment the main strands of Williams's anti-sceptical case, and also to provide its own distinctive style of directly anti-sceptical argument—one that affirms the genealogical primacy of knowledge-permitting contexts. Most generally, I hope to have thereby shown how social epistemology may be fruitfully expanded not only across social space but also across time.<sup>18</sup>

68

## 4 On Saying that Someone Knows: Themes from Craig

Klemens Kappel

In his book *Knowledge and the State of Nature*, Edward Craig (1990) proposes a new methodological approach to one of the crucial questions in epistemology, the question what knowledge is, or more precisely the question regarding the nature and the concept of knowledge. We should, urges Craig, start by asking what we have the concept of knowledge for, what the point or purpose of this concept is. Craig thinks that we can answer this question, and that the answer will illuminate the concept of knowledge. The reason we have the concept of knowledge, according to Craig, is that we use it 'to flag approved sources of information' (Craig 1990: 11). In part, Craig defends his view about the point of the concept of knowledge by providing a detailed *genealogical* account of how the concept of knowledge evolved out of a more primitive concept of the good informant, which in turn arose in a *state of nature*. Moreover, Craig endeavours to show how this throws light on a range of otherwise puzzling features of knowledge and our practice of epistemic appraisals, including perennial issues debated in epistemology, such as scepticism, the Gettier problem, and the controversy over externalism and internalism.

Craig's book strikes one as rich and suggestive. Indeed, it seems very compelling that understanding what we have the concept of knowledge for, and how this concept developed, will help us understand features of the concept of knowledge. And it certainly does seem plausible that this in turn will cast light on the nature of a range of disputes that relate to the question of what knowledge is. Indeed, Craig has many interesting things to say about the point of the concept of knowledge, the development of this concept, and how this explains certain features of our practice of epistemic appraisals, and certain debates in epistemology.

Nonetheless, it is far from straightforward what it is to understand the point of a concept, or how a genealogical story or considerations of a state of nature help us to understand this. How, if at all, does understanding the point and the genealogy of the concept of knowledge contribute to understanding the content of this concept, or the nature of knowledge? On the other hand, one should not dismiss

this sort of project out of hand, and certainly not do so merely on the ground that it deviates somewhat from mainstream epistemology in the analytical tradition.

It is, therefore, worth exploring whether a robust understanding of the kind of approach devised by Craig can be provided. Proposing a distilled version of Craig's approach is the main aim of this chapter. Thus, the issue before us is how to make sense of questions such as 'What is the point of the concept of knowledge?' and 'Why do we attribute knowledge?' Below, I shall characterize and defend a general way of understanding such questions, and contrast it to elements of Craig's approach. On the basis of this, I offer some views about what, in Craig's terminology, is the point of the concept of knowledge,

and again I contrast this to some of the claims put forward by Craig. Like Craig, I believe that understanding the point of the concept of knowledge may yield some insights about knowledge and the concept of knowledge, though space does not permit a discussion of these further questions.

## 1. PRACTICAL EXPLICATION

Often philosophy, particularly in the Anglo-American tradition, proceeds by way of conceptual analysis in some broad sense. This is no less true in epistemology, where much work aims to analyse our concepts of knowledge and justification, and other epistemic concepts.<sup>1</sup>

Clearly, however, there may be other ways of elucidating our concept of knowledge. One apparent way would be by starting with trying to understand the general point of having the concept of knowledge, the function it serves, or the purposes we pursue in making knowledge attributions, or the point of institution of knowledge attributions. This is the approach favoured by Craig, and in his book he offers a rather detailed account of what he takes to be the 'point of' knowledge, backed up by a genealogical story about how the concept evolved.

As a general strategy in philosophy, this is not new, of course. Nietzsche pursued a similar strategy with respect to certain moral concepts, and related themes are found in Wittgenstein. More recently, as Kusch (2009) points out, Kripke, inspired by Wittgenstein, urged us to ask: 'What is the role, and the utility, in our lives of our practices of asserting (or denying) [meaning attributions] under these conditions?' In epistemology, approaches similar to Craig's have been taken up by Bernard Williams (2002) and a number of epistemologists have been directly inspired by Craig's work, and they have to some extent modified or applied Craig's work.<sup>2</sup>

70

These efforts notwithstanding, the kind of approach favoured by Craig has not been widely adopted in recent epistemology. This, I believe, is in part due to the fact that it is often not fully transparent what the approach recommended by Craig really amounts to. Despite a longstanding tradition in philosophy for embarking on just these kinds of enquiries, one may reasonably be somewhat hesitant about our capacity as philosophers to answer questions such as 'What is the point of the concept of knowledge?' or 'How did the concept of knowledge evolve from earlier, more primitive concepts?' The reasons for being cautious are not hard to come by. How does one read off what the point or purpose of a concept is? How does one gain epistemic access to facts about the evolutionary process that lead, or could have lead, to our having a concept with the content of the concept of knowledge?

The aim of this section is to suggest that there is nonetheless a robust interpretation of these 'What is the point of X?-questions', as one may term them. Moreover, it is an interpretation of this type of question that points to a way they can be answered even by armchair philosophers. I shall use the expression 'practical explication' to refer to a

successful answer to 'What is the point of X?-questions' interpreted as outlined below. So, answering the question 'What is the point of *knowledge*?' when the question is interpreted in the way I outline below issues in a practical explication of knowledge.

Note that also Craig uses 'practical explication' to refer to the results of his investigations of the concept of knowledge, though he employs such diverse terms as 'genealogical explanation', 'state of nature explanation', and 'contractarian explanation'. In what follows, however, I shall use 'practical explication' to refer to the kind of explanation I outline here, and not to whatever 'practical explication' refers to in Craig's writings. This act of linguistic appropriation is at least partly justified by the fact that the sound core of the approach proposed by Craig just issues in the kind of practical explication that I outline here, or so I shall suggest at any rate.

Let me first exemplify what I take a practical explication to be. Consider the question 'Why do we have cars?' or 'What is the point of cars?' When read in a very familiar way, the answer is obvious. The reason we have cars is that we have a need for transport, and cars meet that need. More specifically, we have important things to do in different geographical locations, while time and native resources available for transport are scarce. Hence, there is a need for fairly quick and efficient transport. Cars are our chosen means of meeting that need, and that is the reason we have them. So, in this sense we can say that fulfilling a certain need for transport is the point of cars.

Of course, it is correct to point out that people acquire cars for many other reasons as well. For example, people may acquire cars to display their wealth, or to participate in certain life forms. While this is no doubt true, it does not undermine the explanation just offered. There is no reason to deny that items such as cars serve several purposes, though perhaps it is questionable if cars could

71

serve these more symbolic purposes in the way they do, if they didn't serve as means of transport for so many people. At least, it is very unlikely that we would have cars if we never had had the need for transport that they serve to meet.

A practical explication is a kind of explanation, but this explanation is not meant to explain why we have an extensive traffic system based on privately owned vehicles *rather than*, say, a very elaborate system of public transport. Like many other explanations, the practical explication of cars depends on contrast classes. Roughly, saying that we have cars in order to meet a need for transport is meant to explain why we have cars, rather than, say, no means of transport apart from our natural ability to walk and run. It is not meant to explain why we have cars rather than some other equally efficient system of transport.

Having exemplified what I take a practical explication to be, I now want to make three observations about practical explications, all of which should be fairly uncontroversial.

First, obviously, nothing hangs on cars. We can provide similar practical explications of many other things or kinds of things, say social phenomena such as monetary systems,



laws, educational systems, or devices for communication, conventions of etiquette or moral norms. And we can, as is our concern here, try to devise practical explications of various concepts, including the concept of knowledge. Just as we can ask why we have cars, we can ask why we have the concept of knowledge. This is not to insist that any item, or even any artefact, does have a practical explication. Perhaps not every concept has a practical explication of the kind exemplified. Yet, practical explications do make sense, and some concepts do seem to have practical explications, or so I shall suggest at least.

Second, a practical explication of a certain phenomenon proceeds by stipulating first, certain facts about us, such as facts about our physical environment, our biological set-up, our social organization, and second, certain aims or interests that we typically have, or have in certain more specific circumstances. Against the background of such stipulations we can see that there is a certain need, and we can stipulate that the phenomenon to be explained fulfils that need. This is a practical explication of the phenomenon in question.

So, the format of a practical explication is roughly this. To provide a practical explication of some explanandum E, two sorts of claims are advanced:

- (1) Given a set of facts F, and a set of aims or interests I, we have a certain need N.
- (2) E is what actually fulfils the need N.

Note that a practical explication does not attribute intentions or teleological purposes to anyone or anything, apart from what is involved in the attribution of general aims and interests to people. Thus, it is not part of a practical explication that we somehow have the intention to use E in a certain way, though the existence of such intentions may be consistent with a practical

explication of E. And the practical explication does not attribute the teleological property of having a certain purpose. Neither does a practical explication posit a developmental history, or make appeals to some primitive state. It is not part of the practical explication to hold that E developed in certain stages, though it is clearly consistent with a practical explication that it did.

Practical explications appeal to facts, typically facts about the physical world, our biological or psychological constitution, or our social organization. And practical explications invoke facts about our general aims and interests, or the aims and interests we have in more specific situations. Generally, practical explications thus depend on empirical, contingent facts. However, this need not invalidate the use of practical explications in epistemology or in philosophy more generally, since the practical explications we are interested in as philosophers may rely on trivial facts. When this is so, philosophers can appeal to these trivial facts and offer plausible practical explications, even without leaving the armchair, though of course not without running the risk of being wrong even about the trivial facts. Clearly, however, practical explications may also proceed at more detailed levels where no trivial facts can be appealed to, and where

proper empirical investigation therefore is necessary. And in many cases, the practical explications of various phenomena may be hard to come by, for example, because the underlying facts are not easily accessible, even with the best of empirical investigations.

Third, at least in some cases, practical explications elucidate the phenomena explained. In the case of cars, we can understand a range of features of cars by seeing how they relate to the practical explication of why we have cars. Cars have wheels, for example, and can carry a number of passengers, and these features of cars makes obvious sense against the background of the practical explication suggested. Similarly, one might hope to discover features or properties of a concept on the background of a practical explication, or show how such features make sense given a practical explication of the concept.

## 2. PRACTICAL EXPLICATION OF THE CONCEPT OF KNOWLEDGE

Having now what is, I hope, a tolerably clear idea of what a practical explication is, consider the question: why do we have the concept of knowledge? What is the point of the concept of knowledge? We can now rephrase this question as follows: what is the practical explication of the concept of knowledge?

The suggestion I shall elaborate is that we can answer this question by noting a set of features of our cognitive capacities and practical ends, which indicate that we have a certain need, a *conceptual need*, as I will call it. Next, the suggestion is that our concept of knowledge fills the void defined by that conceptual need. As will be clear, it is actually the first step—that of delineating the contours of

73

a particular conceptual need—that is the most important. And it is here that the practical explication is most likely to go wrong. Nonetheless, the proposed practical explication of the concept of knowledge might also fail at the second step: it might correctly identify the shape of a conceptual need, yet the conceptual need is not met by deploying our concept of knowledge, but in some other way. I regard this possibility as unlikely, though not as unthinkable, but I don't offer specific arguments against it.

As one might expect, much of the practical explication of knowledge that I wish to put forward is encapsulated in this quotation from the beginning of Craig's book:

Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. They have 'onboard' sources, eyes, and ears, powers of reasoning, which give them a primary stock of beliefs. It will be highly advantageous to them if they can also tap the primary stocks of their fellows—the tiger that Fred can see and I can't may be after me and not Fred—that is to say, if they act as informants for each other. On any issue, some informants will be better than others, more likely to supply a true belief. (Fred, who is up a tree, is more likely to tell me the truth as to the whereabouts of the tiger than Mabel, who is in the cave.) So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use. The hypothesis I wish to try out is that the

concept of knowledge is one of them. To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information. (Craig 1990: 11)

However, before turning to the main idea suggested in this passage—the idea that the concept of knowledge is ‘used to flag approved sources of information’—I want to consider another important matter. Neither in the passage quoted nor in his book, does Craig consider the special role the concept of knowledge may have for enquiry, independently of whatever need we might have for identifying sources of information. Yet there seems to be a very distinct role for the concept of knowledge in enquiry, or this is what I want to suggest at any rate. Consider what we might call our need for an *enquiry-stopper*. The need for an enquiry-stopper arises out of the following three trivial observations about our interest in truth and certain limitations in our cognitive capacities.

The first is that truth matters. For a range of important propositions we generally benefit prudentially from treating them as true just in case they are true. Here I assume that what makes a proposition important is its relation to our practical or prudential concerns. Treating an important proposition as true, while it is in fact false, is generally worse in prudential terms than treating it as false. And, conversely, failing to treat a relevant proposition as true, when fact it is, is generally less optimal in prudential terms. Some propositions, of course, are not important for our practical

74

concerns, so at best we gain nothing by treating them as true and do not incur the risk of a loss by treating them as false. In certain special cases, propositions are really important for our practical concerns, yet because of unusual circumstances, we are better off not believing a true proposition, or firmly believing a false one. My quality of life may be much enhanced if I falsely believe that the medication will cure my life-threatening disease, or if I do not believe the truth that I have failed in my life's most important project. While there are such cases, it is still true that we generally benefit from treating as true all the important propositions, and this is the claim I wish to make.

Many epistemologists suggest that at least some true propositions matter in the sense that knowing them is in some sense *epistemically* better than merely believing them, or not believing them, where this sense of betterness is entirely unrelated to prudential goodness or prudential concerns. I wish to ignore issues about this sense of epistemic goodness, however, and just rely on the obvious ways in which the truth of propositions may be important for our practical concerns.

The second observation is that without some form of enquiry, we cannot have truth. But enquiry is always costly. At any stage in enquiry, continuing the enquiry comes with certain costs in terms of time and resources. Though facts may sometimes be said to be directly observable, beliefs about facts are nonetheless always acquired by some *method of enquiry* in a particular set of circumstances. A method of enquiry whether *p* may be viewed as a process or mechanism that serves to rule out various possibilities. If the

method delivers the result that  $p$ , then the method does so in virtue of ruling out certain non- $p$  possibilities, and vice versa a method that delivers the result that non- $p$ .

Moreover, we need to worry not only about the error-possibilities that are ruled out by the method we use, but also about those that are not. For any belief, method, and particular set of circumstances, there are undetectable error-possibilities, that is, ways in which things might have gone wrong such that the subject in question would have ended up with a false belief about the relevant subject matter, despite relying on a generally reliable method of enquiry. One's chosen method may malfunction, or there may be unexpected features of the circumstances to prevent it from being reliable under the circumstances at hand. Ruling out error-possibilities by using a method of enquiry and by securing that the conditions for its application applies, generally requires resources in the form of time, energy, and attention. In this sense, enquiry is always costly, and can, from the point of view of the individual, be considered a kind of risky investment.

The third observation is that enquiry, in a sense, has no natural stopping point. Imagination may always reveal further conceivable but as yet uneliminated error-possibilities. Of course, the more resources one has invested into enquiry into some question, the more remote the remaining conceivable error-possibilities may be. Nonetheless, there are always further ways in which things might have gone wrong. As we know all too well in epistemology, it only takes a sceptical mind to realize what such options might consist in. Contrary to what Moore

75

suggested, agents are never in an epistemic position such that they can reasonably infer that any further conceivable error-possibility simply does not obtain. Or, at least it takes considerable philosophical sophistication to discern that one is in a position to rule out any conceivable error-possibility. It is not, of course, that any sound mind would always worry about any remote as yet uneliminated error-possibility. But this is precisely the point that I wish to make. The sound mind will be in need of a way of expressing the attitude that enquiry has now taken one far enough, and that one shouldn't worry about remaining as yet uneliminated error-possibilities.

Perhaps a way of putting the point is this. At any given point in enquiry we face a decision task, which consist in weighing the risk of being wrong about  $p$  against the risk of wasting time and resources on further enquiry whether  $p$  in the event that  $p$  is true. If one's belief that  $p$  is true it would be better just to treat it as true, rather than continue enquiry. But, of course, if one is inclined to think that  $p$  is true, but  $p$  is in fact false, it may be better to spend extra resources in the effort of discovering that it is false. Now, clearly, often it will be very difficult for enquirers to solve this decision problem by calculating the expected utility of the two options. In addition, there will be a point where even beginning to calculate in this way risks being a waste of time and energy, though for obvious reasons it may be even more difficult to decide when this point has been reached.

This suggests that we have a particular conceptual need. As a first approximation, we may describe it as a need to express a certain kind of judgement about the beliefs we acquire during enquiry: that enquiry has now gone on long enough. Or as one might perhaps put it in more detail: in the nature of the case, we always regard the propositional content of our beliefs as true. Yet we often believe a certain proposition  $p$ , and still regard it as a live possibility that  $p$  might be false, despite the evidence in favour of  $p$  that made us believe it. In such cases, there is always a *pro tanto* reason to enquire further into the question whether  $p$ . But often, in such cases, we nonetheless benefit from treating a belief as true, full stop, despite the remaining *pro tanto* reason to continue enquiry. So, we need a kind of device to express the attitude that our epistemic position with respect to a particular proposition  $p$  is now so good that we should stop further enquiry whether  $p$ . We need a way to command a switch of our attention away from further uneliminated non- $p$  possibilities, a way to urge that we simply take the truth of  $p$  for granted in our practical deliberation, as well as in our enquiry into other questions. Or, as I shall say, we have a need for an enquiry-stopper.

In his book, Craig doesn't discuss the role of knowledge in enquiry, and so he does not make claims analogous to what I have said about the need for an enquiry-stopper. Instead, he suggests, as we have seen, that we use the concept of knowledge to identify good informants, and Craig offers a genealogical story about how our concept of knowledge has evolved in response to a need to identify good informants that arose in a state of nature. The basic idea is plausible, I

76

think, but the case for it could be made using the notion of a practical explication that I proposed above, and without resorting to genealogy or state of nature explanations. This is what I shall now do.

Again, the basic approach is to identify a set of facts concerning our cognitive capacities and basic interests, in the light of which we can see that we have a conceptual need of a certain structure. As before, assume that truth matters to us for prudential reasons, and assume that truth is available only through methods of enquiry, the exercise of which are costly, and in some cases not possible. Now, there are two new elements to consider.

First, subjects differ with respect to their cognitive capacities, their decisions about how to employ these capacities, and the circumstances under which they do so. That is, individuals may differ with respect to the cost and possibility of acquiring beliefs with a certain epistemic strength within certain domains, and at particular times they may differ with respect to the propositions they have chosen to adopt as targets for their investigative efforts. This results in a great variability in the strength of epistemic position that various subjects have with respect to many propositions. This variability may be due to differences between subjects regarding how they are located in time and space, and it may be due to variation in internal cognitive capacities and cognitive history, or their ability to exercise these capacities, and their choices about targets of enquiry.

Second, because of this variability, it may be the case that some subject  $S_1$  is not in a position to rule out error-possibilities relevant for some proposition  $p$ , while  $S_1$  nonetheless has evidence that some other subject  $S_2$  is in such position. Or more generally, it may be less costly in terms of time and resources for  $S_1$  to acquire the belief that  $S_2$  is in a certain desirable epistemic position with respect to some target proposition  $p$  than to get in such position herself. Given that such cases are not too rare, it is obviously useful for subjects to be able to identify other subjects who assert  $p$  while being in a position to rule out relevant error-possibilities. Hence, we have a general need to be able to identify properly positioned informants, or good informants, as Craig says, though there are some important differences between Craig's view about what it takes to be a good informant and mine, to which I shall return.

More carefully, what is useful for us to have is a range of predicates that serve to express a variety of epistemic evaluations of the epistemic position of other subjects. If I register that some subject  $S$  believes some proposition  $p$ , I clearly need to assess  $S$ 's epistemic position with respect to  $p$  in order to make any use of that observation in my own cognitive economy. Without this epistemic assessment it is simply indeterminate how the proposition should fit into my own theoretical and practical deliberation. Now, for much the same reasons as I noted above concerning the need for an enquiry-stopper, one useful evaluative epistemic predicate would express a permission (or obligation) to foreclose further enquiry and a permission (or obligation) to rely unconditionally on a proposition in one's practical deliberation. Thus, we have a need for a predicate to express

77

the judgement that we should take the truth of a transmitted proposition for granted.

We can say that it is useful for us to be able to receive or transmit *fully trusted information*, information that  $p$  conjoined with the evaluation that, because of the origin of the information that  $p$ , there is no point in questioning  $p$ , no point in investing resources in further enquiry whether  $p$ , and such that the information that  $p$  should be taken for granted in practical deliberation. Thus, we have a conceptual need to express the judgement that some proposition  $p$  delivered by an informant  $I$  is to be considered fully trusted information.

As before, the need for transmitting fully trusted information, and the conceptual need that this generates, relies on basic and trivial facts about our cognitive capacities, the physical environment, and the different distribution and application of diverse subjects' cognitive capacities. We benefit from transmitting information from one individual to another, and the reason this is a benefit, of course, has to do with the fact that we occupy different epistemic positions, combined with the fact that our own epistemic position is often not as good as we would like. Denying these facts would require holding either that we rarely have an interest in truth, that enquiry comes for free, that cognitive capacities are not unevenly distributed or applied, or that one is never in a position to register that some other subject is better situated than oneself with respect to some subject matter. None of this is plausible, of course.

## 2.1. Outline of a Conceptual Need

I have now indicated the contours of two conceptual needs: the need for an enquiry-stopper and the need for transmitting fully trusted information. We can characterize these conceptual needs in more detail by characterizing the content of the predicate or the judgement that we would need to express the attitude in question. I shall do so for each of the conceptual needs, and then indicate the obvious way in which the two needs may be considered two instances of one slightly more general conceptual need.

Consider first the need for an enquiry-stopper. This is the need for a predicate that permits us to express a judgement (about a subject  $S$  and a proposition  $p$ ) with the following kind of content:

- (1)  $p$ , and  $S$  is in a sufficiently good epistemic position with respect to  $p$ , such that  $S$  ought to take the truth of  $p$  for granted in her practical and theoretical deliberation.

We can add more flavour to what it means to take a proposition for granted by saying that taking a proposition for granted includes that one believes  $p$  with confidence, or perhaps better that one should believe that  $p$  with considerable confidence, that one shouldn't gamble on  $p$  being false, that one shouldn't spend time and resources amassing further evidence for  $p$ , and that one should not spend time looking for evidence against  $p$ , and so on.

78

Consider then the need for transmitting trusted information. What we need is a predicate that in its basic form permits us to express the following type of judgement (about two subjects  $S_1$  and  $S_2$ , and a particular proposition  $p$ ):

- (2)  $p$ , and  $S_1$  is in a sufficiently good epistemic position such that, given the right circumstances of transmission,  $S_2$  ought to take the truth of  $p$  for granted in her practical and theoretical deliberation.

I haven't here said what counts as a right circumstances of transmission; this is obviously an important question, and it is likely to be a very delicate matter, but my intention here is only to describe certain structural features of a conceptual need, not to try to identify the exact circumstances under which information can be transferred as fully trusted from one individual to another.

Clearly, the kind of judgement exemplified in (1) may be viewed as a special case of the judgement exemplified in (2): it is the special case where the informant is the same individual as the recipient. Moreover, the judgement exemplified in (2) is a special case of a more general class of judgements that concern a group of subjects rather than only one as in (2). Hence, while the need for an enquiry-stopper and the need for transmitting trusted information may seem two very different things, they are in obvious respects very similar. Thus these two needs may be viewed as instances of the same slightly more general conceptual need.

Let us introduce the term ‘K-predicate’ as a name for the predicate we would use to express the kind of judgement corresponding to this slightly more general conceptual need. So, the K-predicate would be partly characterized as follows:

$K(S_1, S_1-S_n, p)$  iff  $p$ , and  $S_1$  is in a sufficiently good epistemic position such that  $S_1-S_n$ , given right circumstances of transmission, ought to take the truth of  $p$  for granted in their practical and theoretical deliberation.

In enquiry, the enquiring subject may apply the K-predicate to herself to express the judgement that enquiry has now gone far enough, and in this case right circumstances of transmission almost always obtain. In transmission of information, a subject  $S_1$  may judge that some other subject  $S_2$  is in a strong enough epistemic position with respect to some information that  $p$ , such that  $S_1$  should take the truth of  $p$  for granted in her theoretical and practical deliberation. Or one may use the K-predicate to recommend someone as a source of information regarding a particular subject matter. Thus,  $S_2$  may express the judgement that  $S_1$ 's epistemic position with respect to  $p$  is such that some larger group should take the truth of  $p$  for granted, given the right circumstances of transmission.

The point of introducing the K-predicate is only to provide a convenient way of characterizing and referring to the particular conceptual need that I am interested in. The word ‘need’ should not be understood in the strong sense according to which S has a need for X in order to do F only if S cannot do

79

F without X. In this sense, humans have a need for water, for example, since without water humans die. When I here speak of a conceptual need I have a less dramatic sense of the word ‘need’ in mind; my suggestion is only that, other things being equal, we are better off if we can meet this conceptual need in some way. We are better off meeting the need, though we might also get along without doing so.

It is now convenient to expand the terminology a little. Let's call a state of a subject to which the K-predicate applies a *K-state*, and let us say that the concept thus expressed by the K-predicate is *the concept of a K-state*. I haven't yet said that the conceptual need defined by the K-predicate is actually met by our possession of the concept of a K-state, or that this need could not be met in other ways. Having the K-predicate and the concept of a K-state as a part of one's linguistic and conceptual inventory would meet the need in question, of course, but there might be other predicates and concepts that would address the same conceptual need in a different way. For example, one might imagine rational enquirers who are always willing to consider the possibility that some proposition is in fact false. These would be subjects who never take any proposition for granted in the way I have suggested, but rather attribute some degree of confidence, or say in terms of subjective probabilities approaching but always less than 1.

Neither have I said that we use the concept of knowledge to meet the conceptual need identified by the K-predicate. This would be a further claim about the concept of



knowledge, and as I see it, it is in part a contingent, empirical claim. Finally, I haven't said that the concept expressed by the K-predicate just is the concept of knowledge.

This notwithstanding, it seems clear that the concept of a K-state has a range of features in common with the concept of knowledge, as we would usually understand this concept, and it might be worth briefly outlining some of these parallels. Clearly, just like states of knowledge, K-states are factive, they involve belief conditions, and they require some degree of justification or warrant, however understood. Moreover, being in the K-state entails being subject to certain norms concerning enquiry and practical deliberation: being in the K-state with regards to some proposition  $p$  entails both that one shouldn't continue one's enquiry whether  $p$ , and that the proposition that  $p$  should enter into one's practical deliberation in a very specific way. So, at least we can say that K-states are in many important ways just like states of knowledge.

Moreover, I do believe that our concept of knowledge is what actually serves to meet the need identified, though I haven't argued this further point explicitly. The point of the concept of knowledge is to accomplish the task delineated by the K-predicate. Thus, the role of the concept of knowledge in our cognitive ecology, as Kaplan says, is this: in enquiry, when we believe some proposition  $p$ , we say that we know that  $p$  in order to express the judgement that  $p$  is indeed true, and that our epistemic position with respect to  $p$  is strong enough

80

to merit taking the truth of  $p$  for granted. In transmission of information, we say that some informant  $S$  knows that  $p$  in order to express the judgement that  $S$ 's epistemic position with respect to  $p$  is strong enough to warrant our taking the truth of  $p$  for granted.<sup>3</sup>

### 3. GENEALOGY, STATE OF NATURE, CONCEPTUAL SYNTHESIS, AND THE POINT OF A CONCEPT

Having now outlined what I hope is a workable notion of practical explication, and begun outlining the practical explication of knowledge, it is worth pointing out how this approach differs from that of Craig and related approaches. One important element of Craig's work is that he proposes a genealogy of the concept of knowledge, a genealogy that traces the concept back to a predecessor concept employed in a state of nature.

But what is a genealogy of a concept? Bernard Williams writes:

a genealogy is a narrative that tries to explain a cultural phenomenon by describing a way in which it came about, or could have come about, or might be imagined to have come about. (Williams 2002: 20)

In the words of Martin Kusch, the method proposed by Craig is genealogical, in that it 'seeks to illuminate knowledge by providing a hypothetical historical narrative of the process in which this concept was put together' (Kusch 2009).

A practical explication of knowledge, as I understand it, is not a genealogy in this sense. This is a major advantage, I think, for fairly obvious methodological reasons. Genealogical explanations may, as Williams says, seek to chart the actual history of some item. If so, they require an input of empirical observations that may be very hard to get in the case of the concept of knowledge. It is simply very difficult to know with any confidence how our actual concept of knowledge, or predecessors to this concept, evolved in pre-historic times, or even in more recent times. Of course, we can study the epistemological tradition and learn how philosophers at different times has invested the word 'knowledge' and its cognates with somewhat different meanings, but this may obviously be a very unreliable guide to the relevant fact about the concept of knowledge as it features in our actual conceptual inventory.

The methodological obstacles are no less formidable when we turn to hypothetical genealogies. If we assume, for the sake of argument, that the content or other features of knowledge to some extent reflects the 'hypothetical historical narrative of the process in which this concept was put together', as Kusch says it, then it is obviously quite important to know *which* hypothetical story to adopt. Clearly, there will be many hypothetical stories to choose between, and it is far

81

from obvious how we would identify the relevant hypothetical story among the many possible stories.

By comparison, the practical explication of knowledge that I have proposed relies only on trivial observations of actual facts as they are now; it does not posit a developmental history, and neither is it a hypothetical narrative.

Genealogical explanations sometimes take the form of *state of nature explanations*, that is, explanations that contrast some actual state of affairs with a state of nature, where the state of nature may be an actual historical state, or a hypothetical one. State of nature explanations are, of course, familiar in political philosophy, where they are usually invoked for normative purposes. They usually employ the strategy of showing that some state of affairs—a distribution of resources or whatever—is just (or unjust) in virtue of the fact that it could (or could not) have evolved from the state of nature in the right way. In a related version, a state of nature explanation may seek to offer a story about how we might form a collective intention to adopt a certain convention or social institution, such as the institution of property rights (Hume's example). In other versions the ambition is more like explaining why we have an individual reason to adhere to certain social institutions, such as certain norms of morality (Hobbes's example).

Craig talks about a state of nature, and it even figures in the title of his book. For Craig, however, the point of invoking a state of nature is not to justify any normative stance. The state of nature is a kind of hypothetical primitive human society that contrasts with our actual and presumably more complex and sophisticated social organization, and the

state of nature figures only as part of an explanation of how our concept of knowledge evolved from a more primitive predecessor-concept.

The practical explication of knowledge outlined above is, of course, entirely different from all this. It does not invoke a state of nature, either for normative or explanatory purposes. Again, this is a major methodological advantage for the reasons already rehearsed: we know too little about the actual primitive life forms and the concepts that figured in these life forms. And there are too many hypothetical primitive states to choose between.

In many places Craig writes as if a concept may inherit certain properties or features from its predecessor-concepts. Craig, for example, writes:

The concept of the informant will then leave its mark, so to speak, on the concept of knowledge, and the hypothesis is that for this reason an approach via the former concept will be likely to provide a more illuminating account of the concept of knowledge than will any other method. (Craig 1990: 95)

Indeed, if a concept is not subject to some form of hereditary influence of some sort, it is hard to see how a genealogical explanation can illuminate features of the concept. However, one might reasonably doubt whether there can be the relevant kind of hereditary mechanisms between (tokens of?) concepts, except

82

in some metaphorical sense. In particular, it may be very hard to understand how a purely *hypothetical* narrative can involve the relevant kind of hereditary influence. Why does some story about how a concept *could* have evolved, but did not actually evolve, leave any marks on our actual concepts? To make sense of this, one would have to posit causal mechanisms that cross modal boundaries in the most astonishing way.

At any rate, the practical explication of knowledge put forward above does not assume any process of conceptual inheritance; it is not part of the story that the concept of knowledge carries certain traits that stems from predecessor-concepts. But practical explications are not for that reason bound to be static. Obviously, the conceptual needs generated by facts about the environment, our cognitive capabilities, our general or specific aims and so on, might evolve for the simple reason that the underlying facts change over time.

In various places Craig writes as if the concept of knowledge is ‘constructed’ out of the concept of the good informant, or some similar concept or set of concepts. If taken literally, this idea is rather different from what is no doubt Craig's settled view, namely that the concept of knowledge somehow *evolved* out of a more primitive conceptual device serving to identify good informants. This is not the same as claiming that the concept of knowledge is *constructed* out of the concept of the good informant. Indeed, the view that the concept of knowledge evolved out of something like the concept of the good informant through some evolutionary process is consistent with the view that the concept of knowledge is primitive, that is, not constructed at all out of any more primitive

parts. Anyway, the practical explication that I offer neither affirms nor denies that the concept of knowledge is constructed out of other concepts.

Yet a third idea floated in Craig's book is that we can fairly easily advance plausible hypotheses about the role of the concept of knowledge. As he says:

We take some *prima facie* plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application. (Craig 1990: 2)

And, in a similarly optimistic tone:

One doesn't have to commit oneself to a great deal of epistemology or semantic theory, as the standard approach evidently does, to presume that there is such a thing as the point of this concept [the concept of knowledge], what it does for us, the role it plays in our lives. (Craig 1990: 3)

It has already been made clear, of course, what Craig takes to be a *prima facie* plausible hypothesis concerning the point of the concept of knowledge: 'the concept of knowledge is used to flag approved sources of information' (1990: 11). No doubt, Craig's settled view is that the genealogical story does offer some additional support to his view about the point of the concept of knowledge. Still, it is noteworthy that in these passages Craig advances his views about the point of knowledge without mention of any state of nature or genealogy. He seems to

83

suggest that we can fairly directly see what would be a plausible hypothesis about the point of the concept of knowledge, and then we can move to the question what a concept having this purpose would have to be like.

Craig's assumption that we can fairly easily discern the point of the concept of knowledge is questionable for familiar reasons. In general, we should not assume that the point of a concept is part of its conceptual content, and indeed Craig does not rest his case on this assumption. Hence, while it might be that conceptual analysis can reveal something about the content of the concept of knowledge, it is not really suited for discovering anything about the point of that concept.

It might be tempting to assume that the point of the concept of knowledge can somehow be gleaned from our linguistic practice featuring 'know' or 'knowledge' and their cognates, or by studying the intentions we manifest when using these words in their ordinary meaning. However, there are familiar reasons why this is not plausible. Compare the question: 'Why do we say that things are on fire?' Well, sometimes to scare people, sometimes to encourage people to escape the burning house, sometimes to celebrate (if we are resistance fighters having just succeeded in putting the enemy's property on fire). Something similar may hold for our use of the words 'know' and 'knowledge'. So, strictly speaking, the set of aims manifested when saying that someone knows may be uninformative about the point the concept of knowledge, and this is of

course due to the fact that once we have the concept, we may use it for a rather broad variety of things. The actual intentions manifested in our use of the concept of knowledge may therefore constitute a bewildering array that doesn't permit any simple inference to any view about the point of the concept of knowledge.<sup>4</sup>

This is of course not to deny the possibility that our various uses of the concept of knowledge and the word 'knowledge' are unified by the fact that they all involve assertions of propositions of the form *S knows that p*. Similarly, whether your utterance of 'the house is on fire' is intended to scare someone or to encourage someone to celebrate, this utterance normally asserts the proposition *the house is on fire*. For all I have said, the same may hold for utterances of the form 'S knows that *p*'—they all assert propositions of the form *S knows that p*. It should be quite clear, though, that this type of unity in our use of the concept of knowledge does not permit any interesting conclusion about the point of the concept. It is true, but not informative, to be told that the point of the concept of knowledge is to enable us to assert propositions of the form *S knows that p*.

Finally, there is, of course, a metaphysical problem about what we should even mean by 'the point of a concept', an issue that Craig does not address. We should, for familiar reasons, be quite wary of the assumption that talk of the point of a concept refers to genuine teleological properties of concepts.

84

It is an obvious advantage of the notion of a practical explication proposed here that epistemological and metaphysical difficulties pertaining to talk of the point of a concept are avoided. We can fairly easily see how certain conceptual needs arise in response to trivial facts about our aims, circumstances, and cognitive capacities. And we can explain how we have some epistemic access to these facts.

It is interesting to note that in some places Craig actually characterizes his own approach in very similar terms:

We are attempting a 'state of nature' explanation of a number of facts of conceptual or linguistic practice. Such explanations work by identifying certain human needs and arguing that the practices are a necessary (or at least highly appropriate) response to them. (Craig 1990: 89)

This passage may be read as endorsing precisely what I mean by a practical explication, provided that the passage is read as referring to our actual needs, the needs underlying our actual conceptual practice, not some need we would have in a hypothetical state of nature. On this reading, there is nothing genealogical about the account, and no reference to a state of nature different from our actual state. Be this as it may, I want to suggest that the sort of practical explication of knowledge proposed here is a plausible candidate for a purified version of what Craig actually does in his book, despite his own different characterization of his approach.

#### 4. CRAIG ON THE CONCEPT OF THE GOOD INFORMANT

Let us now turn to Craig's views about the concept of the good informant and the concept of knowledge. As mentioned, Craig thinks that the concept of knowledge somehow evolved out of the concept of the good informant, in part through a process of objectivization, as he says. I shall now add some comments to these parts of Craig's view, mostly to compare it to the practical explication of knowledge that I have suggested. Consider first what Craig says about the concept of the good informant:

We begin by considering it at its most subjective. I am seeking information as to whether or not *p*, and hence want an informant who is satisfactory for my purposes, here and now, with my present beliefs and capacities for receiving information. I am concerned, in other words, that as well as his having the right answer to my question,

- (1) He should be accessible to me here and now.
- (2) He should be recognisable by me as someone likely to be right about *p*.
- (3) He should be as likely to be right about *p* as my concerns require.
- (4) Channels of communication between him and me should open. (Craig 1990: 84–5)

85

This gives an approximate indication of what we should understand by the concept of the good informant, 'at its most subjective', as Craig says (1990: 84). Craig's idea is that this concept of the good informant is a concept belonging to a state of nature, and that our current concept of knowledge evolved from this, in part through a process of 'objectivization', in which certain features of the concept of the good informant are diluted, transformed, or left out:

What remains [of the process of objectivization] will not be identifiable with the concept of a good informant [. . .] but it may still bear certain marks of its origin, and if it does then the concept of the good informant may be the key for drawing them to our attention and making their presence comprehensible to us. It will also bear the marks of the objectivisation. Some of the marks may obscure some of the others. (Craig 1990: 97)

I have three main comments on these parts of Craig's view. First, the process of objectivization is obviously quite crucial in these and similar passages. Objectivization is simply the backbone of the genealogical process that takes us from the state of nature to our present situation. However, it is not easy to get a firm grasp of the process that Craig refers to. The driving assumption seems to be that objectivization is a response to certain pragmatic pressures that arise. For example, Craig asserts that 'as objectivisation proceeds', the requirement of 'the detectability of the property will be diluted' (1990: 90). So, the idea is that while the primitive or most subjective version of the concept of the good informant includes a requirement of 'detectability', as is stated in (2) above, this requirement is abandoned as the concept evolves towards our mature concept of knowledge. This is why our concept of knowledge doesn't include any such requirement:

someone S can know that  $p$  though no detectable property of S correlates with S's reliably believing that  $p$ , just as S can know that  $p$  despite being unwilling to tell.

Craig here trades on assumptions about how our concepts would develop in response to certain pragmatic pressures. But it is obviously quite hard to assess the plausibility of this line of thought, and consequently an explanation of features of the concept of knowledge depending on views about objectivization is bound to be very fragile. By comparison, the practical explication of knowledge that I have suggested does not depend on anything like a process of objectivization.

Second, note that perhaps the most striking feature of Craig's way of characterizing the concept of the good informant is that a range of features necessary for identifying and making use of others as informants is built into the very concept of a good informant, viz. the conditions listed in (1) to (4). So, the concept of the good informant comes with a range of epistemic and pragmatic features that enables us to identify individuals in possession of the right answer, and enables us to utilize this capacity of these individuals. It is far from obvious why such epistemic and pragmatic features would have to be built into the concept of the good informant itself. According to Craig, we need the concept of the good informant because, as enquirers, we benefit from being able to identify

86

individuals whom might have the right answer to our question. Nonetheless, the concept of the good informant is not simply the concept of someone having the right answer to some question, according to Craig. But why not? Why not simply stick to the simpler idea that we need the concept of someone being in a good enough epistemic position with respect to some proposition, and then letting the question of the detection of this property, and other questions pertaining to the use we can make of this, be further questions? It is not at all clear why these epistemic and pragmatic features would have to be built into the concept we need. What Craig suggests about the concept of the good informant is rather like the claim that the concept of a good car is the concept of a car that I own and which is parked nearby, and where the keys are presently in my pocket. Clearly, this is odd.

One should note, of course, that Craig is not offering an analysis of our *actual* concept of a good informant. Rather, his concern is to identify a concept that we possessed in some state of nature, and which later evolved into our concept of knowledge. So, while Craig is suggesting a concept that we would need in some state of nature, our discomfort with Craig's notion of the good informant may at least in part be due to the fact that this notion fails to match our current concept of the good informant. Clearly, any discomfort in having this source might not be well placed. Nonetheless, we are still in need of an explanation of why the state of nature concept of the good informant has the odd features that Craig assumes. In comparison, the practical explication proposed above does not in similar ways rest on assumptions that *prima facie* seem peculiar.

Third, the issues just raised bear an interesting similarity to an objection that Craig responds to. This is what he calls the cart before the horse objection: 'Shouldn't being a

good informant be explained in terms of knowing whether  $p$ , rather than the other way around?' (Craig 1990: 94). Or, as Craig states the assumption behind the objection: 'knowing seems to be that central state of the subject which fits him to be a good informant, other circumstances concurring' (1990: 93).

On the one hand, it appears obvious why one would consider something like the cart before the horse objection. Clearly, we should worry that our actual concept of the good informant is something like the concept of someone who knows. On the other hand, as I have just argued, Craig's concern is not to propose an analysis of our actual concept of the good informant. Neither is his concern, one can add, to propose an analysis of knowledge in terms of the concept of the good informant. Rather, his concern is to identify a particular state of nature concept that later developed into the concept of knowledge. Against this background, it is actually not so clear what the point of the cart before the horse objection would be, or why Craig even devotes it any attention.

It is even more striking to note Craig's considered response to the cart before the horse objection:

knowing is best understood in terms of the idea of a potential informant, that we get at its nature by first considering the business of being a good informant and then subtracting

87

something from the result, or modifying it in some way. The concept of the informant will then leave its mark, so to speak, on the concept of knowledge, and the hypothesis is that for this reason an approach via the former concept will be likely to provide a more illuminating account of the concept of knowledge than will any other method.

(Craig 1990: 95)

Now, what Craig seems to be saying here, at least in the most obvious and straightforward interpretation, is that the concept of knowledge is best understood with reference to the *role* it plays in our notion of the good informant. What is striking is that this is very different from saying that the concept *evolved* out of the concept of the good informant through a process of objectivization. Clearly, one can hold that the concept of knowledge is partly shaped by the role this concept plays in our efforts to identify good informants, and yet deny that the concept of knowledge evolved out of the concept of the good informant.

What we have seen, then, is that there really are two very different themes in play in Craig's discussion. One is the idea that the concept of knowledge evolved through a genealogical process, the process of objectivization of the concept of the good informant. The other and entirely different idea in Craig's thinking about these issues is that we might think of the concept of knowledge as being partly shaped by the role this concept has in the concept of the good informant. This latter idea is plausible, I think. It seems to be just an instance of the sort of practical explication of knowledge that I have proposed.<sup>5</sup>

88



## 5 The Swamping Problem Redux: Pith and Gist

Jonathan Kvanvig

The Swamping Problem is one of the central problems in the new value-driven approach to epistemology. Issues concerning epistemic value, however, are not new. We can find them first in Plato's dialogue *Meno*, where Socrates and Meno have a discussion about what type of guide one should prefer if one wants to get to Larissa. The first suggestion is that one should want a guide who knows the way, but Socrates notes that a guide with true opinions will work just as well. Meno's response to Socrates' objection raises precisely the correct worries: he wonders about the nature of knowledge (whether it differs from true opinion) and about the value of knowledge (why we prize it over true opinion). The point of view of value-driven epistemology is that a myopic focus on the former question is a mistake and that questions about the value of knowledge and other epistemic phenomena are important in their own right. Addressing them explicitly can help evaluate the adequacy of proposals concerning the nature of knowledge and other epistemic phenomena.

The natural suggestion in response to Socrates' objection is to look for something valuable to add to true belief so that the composite is more valuable than mere true belief and where the composite is identical to knowledge. Here is one such suggestion: knowledge is justified true belief. This suggestion is appropriate since justification is itself valuable. Sober reflection reveals that one should prefer justified beliefs to unjustified beliefs. So justified true belief should be more valuable than true belief alone, and thus if knowledge is justified true belief, we can explain to Socrates and Meno why we prefer guides who know the way to Larissa: not because they are better at getting us there, but because knowers are in a preferable state to mere true believers.

Sometimes simple explanations arouse legitimate suspicion and such is the case here. Note first that the above explanation may rely on a principle of value additivity: take anything valuable, add something of value to it, and you get something even more valuable. Yet, this principle is false: if a little salt adds flavour and hence value, multiple iterations do not continue to add value.

It is this failure of additivity, I will argue, that it is the pith and gist of the Swamping Problem. This account of the problem has not been properly

appreciated, as can be seen by examining some recent proposals for escaping the problem. We will turn to such proposals after first clarifying what we are trying to explain when we attempt to explain the value of knowledge, and how the issue of additivity is central to such explanations. In particular, we will see that though a social approach to the value of knowledge has a significant role to play in addressing the value problem, its role is partial and incomplete, requiring supplementation by some more general solution to the Swamping Problem.

## TASK 1: WHAT NEEDS TO BE EXPLAINED

One way to proceed here is to notice different levels of suspicion concerning the value of knowledge, beginning with fully general suspicion that knowledge has little or no value at all. One may respond to such suspicion either in terms of the nature of knowledge or in terms of its function in ordinary life. On the former score, one might point out aspects of knowledge, such as true belief, that are valuable, and defend the value of knowledge indirectly in this way. Such defences, however, are likely to have been anticipated by the iconoclasts in question, since the relationship between knowledge and true belief is such an obvious one. Perhaps a more promising strategy is through appeal to the function of knowledge in ordinary life.<sup>1</sup> Such an explanation of the value of knowledge will almost certainly be a social explanation. For example, it is of central importance for us to be able to keep track of information and sources of information that can be trusted, and one of the functions of attributing knowledge to sources is for such tracking purposes.

Such an explanation latches onto something important and significant about the social function of the concept of knowledge. Whatever else we might wish to say about the value of knowledge, there is no question that a good part of the value of knowledge needs to be explained in precisely this way. We can gain a great deal of insight into the value of knowledge by attending to the functional role the concept of knowledge plays in our lives, and we might even try the tack of attempting to understand knowledge, not in terms of necessary and sufficient conditions for the application of some concept expressed by ordinary language, but rather in terms of whatever plays the knowledge-role in ordinary thinking. Initial suggestions along these lines have played an important role in recent epistemology, as epistemologists have suggested that knowledge is necessary for appropriate assertion and sufficient for practical reasoning.<sup>2</sup> The former principle

90

is already social in nature because of the social function of assertion, and the role of knowledge in practical reasoning should be conceived in social terms as well if we want the principle to state both necessary and sufficient conditions. If I know that my sources of information know of what they speak, that is perhaps ordinarily sufficient for me to know as well, but my legitimate reliance on that information in practical reasoning is not determined by my own epistemic state in such cases. If you know that the bridge is safe, but for some reason I don't know that the bridge is safe, it can still be acceptable for me to cross it on your word in spite of my lack of knowledge. The epistemic division of labour involved in trusting engineers, doctors, lawyers, priests, pastors, and social policy gurus is not held hostage to the vagaries of the precise epistemic state of the one who trusts. Individual knowledge, for example, is undermined by exaggerated doubts and by subjective defeaters, but reliance on the knowledge of the authorities in question is not automatically undermined by the presence of such idiosyncrasies.

Such a theoretical approach to epistemology might go farther, attempting not only to account for the value of knowledge in terms of such guiding principles, but also claiming that the resulting theory is all the theory of knowledge we need. Such an extension would

count as a Lewisian approach to epistemology, wherein one's concern is not so much with the nature of knowledge itself but with whatever plays the knowledge role in ordinary life and thought, much as Lewis attempts to understand the role of the idea of chance in ordinary thought in terms of the Principal Principle (and its successors).<sup>3</sup> I will not here explore this project, but it is a worthy one indeed, one fully deserve of the title of being a Lewisian theory of knowledge, one designed from the beginning as a social theory of knowledge that aims first at accounting for the value of knowledge by focusing on its function in ordinary life and thought. Such a theory would be a substantive and original contribution to contemporary epistemology, even though it would likely be at odds with the only extant Lewisian theory elicitable from his idea that knowledge is a quite elusive thing (since the more elusive and contextual a theory goes in epistemology, the less important knowledge seems to become).<sup>4</sup> Whatever the guiding principles of the role that the concept of knowledge plays in everyday life, the social function is certain to be central to that role. In such a theory, then, the social function that knowledge plays would be highlighted and the resulting theory would be able to explain the value of knowledge at least in part in terms of this social role.

It is important to note, however, that whatever the merits of a social epistemological account of the value of knowledge, such an account has no capacity to solve or address the Swamping Problem, absent some really excessive

Procrustean manoeuvres. The Swamping Problem attacks a very special and unique kind of value that knowledge seems to have, that involved in being preferable to true belief and any other combination of its proper parts. Whatever role is played by the notion of knowledge in ordinary life is not likely to be sensitive to this feature, unless we claim that it is part of the role from the start. Such a claim has much to be said against it, however. Whatever role knowledge plays in ordinary life is more likely to involve the value of knowledge rather than the special or unique value it has over its proper parts. A more fitting explanation here is that the factors involved in the role in question are more coarse-grained than that, and that the fine-grained special and unique value in question will be tramped under foot in the process of providing an account of the role, rather than the nature, of knowledge. In slogan form, such a role, or functional, theory of knowledge is better suited to providing an account of the value of knowledge than it is to providing an account of the special value knowledge has over its subparts.

Such a conclusion fits well with some rather sceptical thoughts John Greco (2009) has voiced about the special value in question. Greco doubts that ordinary thinking ascribes any such special value to knowledge, but should be understood instead to involve the more simple attitude that knowledge is valuable or that knowledge matters. I agree with this sentiment: the issues that threaten the special value of knowledge do not threaten the value of knowledge, and the assumptions that give rise to the problem concerning the special value of knowledge are not present in any direct or immediate way in ordinary thought about knowledge. Assumptions and presuppositions of thought are too far below

the surface to sustain any such general claim, and may not be mental states at all. But what follows from this point is that a theory of the value of knowledge designed for use in a functional theory of the sort indicated is not likely to need any account of the special value of knowledge, whereas a theory that attempts to say what the nature of knowledge itself is will require such an account.

This point raises important issues in meta-epistemology and meta-philosophy more generally, and the functional role theorists will insist that anything beyond a functional role theory would be suspect from the start. I do not believe this for a moment, and view such attitudes as excessively procrustean. A focus on this issue will take us too far afield, but it is worth mentioning here since everything that follows is predicated on this point of view. In what follows the guiding principle will be about the nature of knowledge itself, not some substitute designed to play the role that knowledge actually plays, and the presupposition of such a project is that reality cannot be delimited in the way imagined by theories that attend only to roles in ordinary thinking.

Whether or not I'm right about such meta-philosophical issues, the point remains that a large part of value-driven epistemology will involve the social aspects central to such a functional account. But the nature of knowledge itself may remain untouched by such an account, and if so, the special value

92

of knowledge may remain unexplained. For the special value is discovered only by comparing the value of knowledge with the value of its parts, and such a comparison plays no role in the functional approach. The point that bears emphasis, then, is this: the Swamping Problem is, at its core, a problem concerning the special and unique value of knowledge over that of its parts, rather than a problem concerning the value of knowledge itself.

In order to see how the Swamping Problem threatens this special and unique value, we need a clearer understanding of the scope of this value, and here some distinctions are needed to see precisely what the Meno problem is. We can begin by noting what the idea that knowledge is more valuable than true belief doesn't involve. It doesn't involve the claim that all things considered, it is always better to know than to merely believe truly. Suppose my entire family will be killed immediately if it is discovered that I know, rather than merely truly believe, that they are in Dallas today. Then, all things considered, it is not better to know. But such a conclusion shouldn't be thought to show anything about the universality or necessity of the value of knowledge. What it shows is that we shouldn't think of the value of knowledge in terms of all-things-considered value. Instead, the obvious reply is to distinguish between *prima facie* and *ultima facie* value.

We might put this point by saying, rather uninformatively, that all else being equal, knowledge is more valuable than true belief. For some approaches to value, this distinction between *prima facie* and *ultima facie* can be put in the language of defeasible reasoning. On these approaches, when reasoning about preferences and values in general,

finding out that a given state is a state of knowledge is a defeasible reason for preferring that state or concluding that the state is all-things-considered valuable, and finding out that two states differ in that the first is a state of knowledge and the second of mere true belief is a defeasible reason to prefer the first to the second and to conclude that the first has more value than the second.

On some more particularist understandings of value, the value of the parts will depend on the value of the whole, so that knowledge will have *prima facie* value in this way only when it is part of a whole which is *ultima facie* valuable. On such approaches,<sup>5</sup> it is false that knowledge always and necessarily has more value than mere true belief, even when the kind of value in question is *prima facie* value. But even these approaches to value had better not deny a global supervenience thesis about value, to wit, that there cannot be two worlds that differ at the value level without differing in some way at the non-value level. Given the global supervenience thesis, if we find a case where knowledge is not more valuable than true belief, there will be a reason we can cite that is independent of the mere fact that the state in question is a state of knowledge rather than a state of mere true belief to explain why knowledge isn't more valuable in that case, and this

93

explanation will give us a factor that we can control for and hold fixed in order to explain the nature of the claim that knowledge is more valuable than mere true belief.

It is also important to distinguish the question of the value of knowledge from the related issue of our valuing of it. Valuing is things done by persons, and it is a thoroughly prosaic fact to note that people sometimes value things that have no value at all. When we ask questions about the value of knowledge, we are asking about matters that go beyond individual preferences and valuing. Though there may be some ultimate connection between objective value in this sense and our subjective valuing, the connection will not allow us directly and immediately to read off claims about the value of knowledge from claims about the way in which we value it.

Those convinced that objective value can only find a place in the natural world by being definable in terms of subjective valuing will be tempted to specify counterfactual conditions for the identity of value with valuing. For example, one might think that value could be identified with valuing by fully rational individuals apprised of all relevant facts. Such counterfactual accounts find grave difficulty avoiding the conditional fallacy,<sup>6</sup> and give way quickly in the face of these problems to 'better self' theories, theories that identify objective value with what our better selves would view as appropriate valuing on our part.<sup>7</sup> Whatever the connection posited, however, the important point to note is that the kind of value in question concerning the value of knowledge is a kind of value that requires the same treatment as other objective values. The value of human life, for example, cannot be explained sufficiently in terms of what a fully informed individual *might* value, or what our better selves *might* attach significance to; instead, the only hope for such accounts would be to refine the idea that the value of human life involves the valuing a fully informed individual or a better self *would* make.

Anything weaker would have no hope of explaining the value in question. Just so, the value of knowledge cannot be explained either in terms of value judgements a well-placed individual *might* make, but must rather advert to valuing such an individual *would* make. In the light of the difficulties presented by the conditional fallacy, it is doubtful that such a theory can be successful, but we need not enter that fray here, since the only point needed for present purposes is one concerning the kind of strength such an account must employ to have any hope of success. The special and unique value of knowledge is thus a special truth-related, *prima facie* value that it possesses necessarily if at all, and it is a value explicable in terms of valuing only to the extent that the individual doing the valuing is idealized and the valuing in question is a valuing that such an idealized individual would do rather than simply might do.

94

This way of putting the point allows a nice analogy with the theory of probabilistic causation. An initial idea concerning probabilistic causes is to think of them in terms of probability enhancers, but such a proposal founders on the fact that causal factors often interact. Taking birth control pills causes thrombosis even though women taking such pills decrease their chances of getting thrombosis (because they decrease their chances of getting pregnant which also causes thrombosis). Just as a causal factor can be present but masked by other factors, so can the value of knowledge be present but masked. And just as we should expect to uncover probabilistic causes by finding probabilistic enhancement in circumstances in which we have the right controls for interference in place, just so we should expect to uncover the value of knowledge when we impose the relevant controls here as well. So, to claim that knowledge has special value only sometimes and in some places is not to account for its special value, but rather would be like resting content with the claim that causation sometimes involves probabilistic enhancement and sometimes doesn't when faced with counterexamples involving interaction to the persuasive initial suggestion that probabilistic causes are probability enhancers. The right response to the initial counterexamples is a refining of the initial idea, not a revisionary conception that remains complacently comfortable with the opinion that there is no deeper explanation than that the original account is true sometimes and in some places. The hard task of stating precisely the controls that are needed I have not done here, but the idea of doing so is just as plausible here as it is with the notion of probabilistic causation.

## TASK 2: THE TWO-TIERED VERSION OF THE SWAMPING PROBLEM

Given this understanding of the claim that knowledge has special value in terms of the idea of a defeasible value that is always and necessarily present, we are now in a position to state more clearly the nature of the Swamping Problem for this special value. The problem, at bottom, arises when a theory tries to explain the value of knowledge over that of true belief by citing a property that is not always and everywhere additive in value. Once we notice that true belief is valuable, and one identifies knowledge with a property that does not always add value to true belief, we smell the scent of a problem. Of course, the problem can disappear if the cases where value is not added can be explained away in

terms of the other kinds of things that must be controlled for in order to reveal the special value of knowledge, as noted above. But such explaining away offers little hope for the usual kinds of theories the Swamping Problem plagues. What is needed is to avoid *non-additive* properties: properties that, once the relevant controls are in place, do not add to the value of true belief. One of the most obvious of such properties is the property of statistical likelihood of truth.

My preferred formulation of the Swamping Problem employs this property and proceeds in two steps. It begins with an example in which it is clear that such a property adds no value whatsoever, and then treats the example as an instructive analogy about one particular epistemological theory that clearly succumbs to the Swamping Problem. After seeing what a theory cannot be like if it is able to avoid this problem, a challenge can be issued to the reliabilist (as well as other theories) to explain what value-adding feature their theory possesses that is lacking in the theory that succumbs to the problem.

First, then, the example. Suppose I want chocolate. I Google to find places close to me. I get two web pages: one titled ‘places that sell chocolate in Waco’; the other ‘places likely to sell chocolate in Waco’. We may assume accuracy for both lists, and that the second list is generated from correlations: places that sell food are likely to sell chocolate, places that sell sweets are too. We do not assume, for purposes of the example, that the accuracy in question is known to obtain, though we do assume that no special doubts are present regarding accuracy. We assume, simply, that Googling gives us the means to achieve the goal of getting chocolate, one a list of probabilistic means and the other a list of means *simpliciter*. We also assume that there is no evaluative dimension to our appraisal of the lists other than in terms of the goal of getting chocolate: categorization of means in terms of better or worse is restricted to their capacity to serve the goal of getting chocolate. We exclude all other values from consideration, values such as having fancier names for a place of business, finding a business run by an acquaintance, proximity to my present location, and so on.

So I have two lists relevant to the goal in question. We then note four important facts about these lists. The first fact is that the second list is better than nothing, given that the only relevant goal is the goal of getting chocolate. If I want chocolate, rather than choosing randomly from all the businesses in Waco, it would be better to select randomly from the list of places likely to sell chocolate in Waco. The second fact is that, for the purpose in question, the first list is more valuable than the second. I do not wish to say that it would be irrational to rely on the second list, given only the purpose of getting chocolate, but it would certainly be better to rely on the first list. The third and fourth facts arise from introducing a third list, which is the intersection of the first two lists: places that both sell chocolate and are likely to. The third fact is that the intersection list is also preferable to the second list: it would be better to rely on a list that included places that both sell chocolate and are likely to than to rely on a list that included only places likely to sell chocolate. The final fact, however, is the swamping fact: the intersection list

is no more valuable than the first list, given the interest in question. If all I care about is chocolate, it would be no better to use the list of places that both sell chocolate and are likely to than to use the list of places that sell chocolate. Choosing a place of business from the intersection list may be no worse for purposes of getting chocolate than choosing from the first list, but it is certainly no better. Someone who adopts the practice

96

of selecting from the intersection list rather than the first list does not adopt a *superior* means of securing the goal in question, given the assumption that the first list is an accurate list.

One may start to follow red herrings of decision theory to try to avoid this conclusion by claiming that it is not enough to assume that the first list is accurate. Instead, it might be claimed, one will need to assume that one believes with certainty, or knows, that the list is accurate. The second claim is surely false: one can reasonably act on the basis of assumptions that do not constitute knowledge. It is plausible to think that the same point holds for the claim that subjective certainty is needed: one can reasonably act on the basis of full belief with no need to adopt further strategies to take into account the possibility or chance of error, even if one resists subjective certainty about one's beliefs. Some chances of error need not be taken into account when acting, and the case as imagined is of this sort. What we have is a low-stakes situation in which the loss of securing the goal is not important enough to get anxious about the chance of failure. So, the lists are taken at face-value, assumed to be accurate in what they represent, and a choice of means to getting chocolate ensues. In such a case, the intrusion by a friend who insists on constructing an intersection list and counsels its use over the first list would be a case of inapt advice. By my lights, it doesn't cut my chances of missing out on chocolate and I have no interest in practising set-theoretic operations for their own sake. I just want chocolate, and the intersection list does not present any improvement for purposes of getting chocolate over the first list.

I belabour the obvious here because there will be a temptation to niggle about the example once one sees its implications. For those still inclined in this direction when the consequences are observed, I recommend a small experiment: present the case as I did originally to undergraduate classes, and see what response is elicited. I predict your experience will match mine: students think the idea of constructing an intersection list and claiming that it is a better list to use for purposes of getting chocolate is *just crazy*.

What I do not claim, it should be noted, is that it is impossible for a person to value the intersection list over the first list. Someone could begin to worry about the accuracy of the Google search engine, about the methods used to generate the lists, and could rationally judge that more research is required before embarking on a trip to get chocolate. Such a person could feverishly seek to calm the anxieties present, and view finding many different web pages mentioning a place of business together with chocolate as a sign that the chance of error was being reduced of losing out on chocolate whenever



intersection lists are constructed in this way. The more web pages from which to create intersection lists, the better the resulting intersection list, so our anxious friend reasons.

My claim is not that it is impossible to imagine an individual who values the intersection list over the first list, or that it is impossible for such valuing to be rational. My inclination is to think that such individuals are making a mistake

97

of a certain sort: they show too much anxiety about the chance of missing out on achieving a rather low-stakes goal. But even if they are not making a mistake, even if their valuing of the intersection list is rational, it is still optional. It is at most a kind of valuing of the intersection list that a fully informed person *might* make, but it is certainly not true that it is a valuing that any fully informed person *would* make or *needs* to make. In such cases, it is a mistake to take the valuing in question as implying any objective value of the intersection list over the first list, as we noted earlier when discussing the relationship between values and valuing. The intuitive point thus stands that the intersection list is no more valuable than the first list for the purpose in question.

This last fact is the first tier in the structure of the Swamping Problem. The second tier is the analogy between the arena of action and the arena of belief. Instead of thinking in terms of actions based on lists, we think in terms of beliefs that correspond with items in the book of truth and in the book of claims likely to be true. Just as actions corresponding to items on the first list (the list of places that sell chocolate) are preferable to random actions, so are beliefs corresponding to items in the book of truth preferable to random beliefs. Moreover, just as actions corresponding to items on the second list (the list of places likely to sell chocolate) are preferable to those that correspond to items not on the second list, so are beliefs that correspond to items in the book of claims likely to be true as compared with beliefs that correspond to items not in this book. Moreover, just as in the case of action, beliefs that correspond with items in the book of truth are preferable to those corresponding with items in the book of claims likely to be true. Furthermore, if we create a new book which is the intersection of the two books already mentioned, we get a further fact, the fact that a belief corresponding to items in the intersection book is preferable to ones that correspond to items in second book, the book of claims likely to be true. Finally, we also get the swamping fact: there is no advantage to having a belief that corresponds to some item in the intersection book over having a belief that corresponds to some item in the book of truth. In short, truth plus likelihood of truth is not preferable to truth alone.

The analogy thus presents an argument for four claims about epistemic value, if our theory of justification identifies it with statistical likelihood of truth: first, true belief is better than false belief; second, justified belief is better than unjustified belief; third, justified true belief is better than justified belief; but fourth, and most important in our context, justified true belief is not better than true belief. These claims about preferability and value are, of course, claims about epistemic preferability and epistemic value,

specifiable in terms of the epistemic goal as usually formulated, in terms of getting to the truth and avoiding error.

This much of the Swamping Problem is, I think, untouchable. It is untouchable in the initial case that adding statistical likelihood alone doesn't add value, and

98

it is untouchable that the analogy shows that one had better not identify justification with statistical likelihood of truth. It is also clear that substituting some other notion of objective probability for the statistical notion won't be much help either. For example, suppose one's theory of objective probability is a propensity theory. It still remains true that a list of places that sell chocolate is better than a list of places with a propensity to sell chocolate, and that there is no advantage relative to the pursuit of chocolate in favour of a list of places that both sell chocolate and have a propensity to do so over a list of places that sell chocolate. The lesson here is that objective probability alone had better not define one's theory of justification if one hopes to solve the Swamping Problem. Hardly any theories of justification, and certainly none of the major players in the contemporary scene, rely on such an identification, however, so the question arises how we are to extend this result to raise any problem for ordinary accounts of justification. In typical presentations of the Swamping Problem, the target is externalist theories of justification, most especially the quite popular process reliabilist version of externalism. So how do the results achieved so far cause a problem for process reliabilism?

I have argued with others that there are versions of externalism that have such resources but that process reliabilism does not cite features with such resources.<sup>8</sup> After all, if objective probability itself succumbs to the Swamping Problem, why would the fact that there is an etiological relationship to a process or method responsible for that probability relieve the theory of the problem? Such a causal relationship to methods or processes does not seem to be the kind of feature that adds value beyond the value of true belief, so there is no apparent reason here to think that ordinary process reliabilism is in a better condition with respect to the Swamping Problem than is the simple objective probability theory. Resistance to this claim can be found in recent literature on the subject, but I will argue that these defences of process reliabilism fail to come to grips with the real Swamping Problem as described above, and instead only succeed at blocking certain other versions of the problem that are much easier to avoid.

### TASK 3: THE PROPOSALS OF GOLDMAN AND OLSSON

In a recent article, Alvin Goldman and Erik Olsson (2009) propose two different ways to try to respond to the Swamping Problem. Both approaches go beyond

99

the identification of justification with objective likelihood of truth, and thus provide some hope of avoiding the Swamping Problem. The first approach is the conditional probability approach:

What is this extra valuable property that distinguishes knowledge from true belief? It is the property of *making it likely* that one's future beliefs of a similar kind will also be true. More precisely, under reliabilism, the probability of having more true belief (of a similar kind) in the future is greater conditional on S's *knowing* that  $p$  than conditional on S's *merely truly believing* that  $p$ . Let's call this proposed solution to the EVOK problem *the conditional probability* solution. Probability should here be interpreted objectively.

(Goldman and Olsson 2009: 28)

The claim here is that if we condition on knowing versus truly believing, the likelihood of being correct in the future goes up. But obviously not for just any old kind of belief, but for a special kind. Which kind is that? Well, for 'beliefs of a similar kind,' the probability goes up when conditioning on knowledge rather than on mere true belief.

Once we appreciate the nature of the Swamping Problem as a problem concerning properties of belief that are non-additive of value in the presence of true belief, it becomes hard to see how the above proposal is helpful at all. In the analogy involving chocolate, I don't even know how to begin thinking about applying this idea to new businesses of the same type, conditional on the first list (places that sell chocolate) and the third list (places that both sell chocolate and are likely to).

There is an answer to my perplexity, however. The proposal is not really the proposal above involving conditioning only on the two lists in question, but a rather more complex proposal involving several contingencies. Goldman and Olsson note that the proposal depends on contingent regularities in order to derive the conclusion that one of the conditional probabilities is higher than the other, and they say,

To see what roles these regularities play, suppose S knows that  $p$ . By the reliabilist definition of knowledge, there is a reliable method M that was invoked by S so as to produce S's belief that  $p$ . By non-uniqueness, it is likely that the same type of problem will arise again for S in the future. By cross-temporal access, the method M is likely to be available to S when this happens. By the learning assumption, S is likely to make use of M again on that occasion. By generality, M is likely to be reliable for solving that similar future problem as well. Since M is reliable, this new application of M is likely to result in a new true belief. Thus the fact that S has knowledge on a given occasion makes it to some extent likely that S will acquire further true beliefs in the future. The degree to which S's knowledge has this value depends on how likely it is that this will happen. This, in turn, depends on the degree to which the assumptions of non-uniqueness, cross-temporal access, learning and generality are satisfied in a given case.

(Goldman and Olsson 2009: 29–30)

So the claim is not that the conditional inequality explains the value of knowledge over that of true belief. The claim is, rather, that when certain contingent

features are in place, we should expect the conditional probability of future true beliefs to be higher, given that one knows rather than that one merely truly believes.

This response, however, fails to come to grips with the core of the Swamping Problem. As we have seen, once the relevant controls are in place, we should expect, always and necessarily, for knowledge to be a value-enhancing characteristic of a state of true belief, not just such a characteristic when the person is likely to use the same method in the future and the world hasn't changed so that what was reliable in the past is no longer so. To uncover the special value of knowledge, we have to control for interaction by values outside the purely cognitive sphere, but once we do so, we should find that knowledge is special. It is not only special when the future resembles the past and when people retain their dispositions across time of how to find out what the world is like. Perhaps we are stuck having to adopt such a revisionary account of the value of knowledge, but we shouldn't offer such an account and pretend that it is not radically revisionary.

Why do I insist that this account is revisionary? I spent some time on this issue in my book,<sup>9</sup> and here I will offer a brief reminder. There are worlds in which my coming to know one more thing makes certain that the world will end. So suppose I just now have come to know one more thing. In such worlds, the likelihood of my having any future beliefs at all, true or otherwise, conditional on knowing versus truly believing, is the same: zero. But this probability does not undermine the fact that I ended my cognitive career on a high note. I *learned* something, and that is a good thing. There may be other more important things on which to end a life: bringing about world peace, for example. And certainly, from an all-things-considered perspective, it would have been better had I not achieved this additional learning. But merely acquiring a true belief wouldn't have been as good from a purely cognitive point of view, and we deserve an explanation of that. In this way, and in others discussed in my book, the conditional probability of the future offers no hope of being able to explain the special value of knowledge.

Goldman and Olsson have another proposal. It involves a main idea and what they term a 'subsidiary' part. The main idea is that reliable processes come to possess 'autonomous' value, value that is not dependent on the value of the resultant true beliefs. The subsidiary part arises because they believe that the Swamping Problem 'assumes that, according to reliabilism, the value of a token reliable process derives from the value of the token belief it produces' (2009: 31) and they thus wish to argue that a token of a reliable process has value even when it does not produce a token true belief. I think they are mistaken in this concern. Nothing in the account of the problem described above assumes that reliable process tokens have value only when they produce true beliefs. Perhaps there is a need for an explanation as to why reliable false beliefs are more valuable than

unreliable false beliefs, but a defender of the Swamping Problem need hold no brief on this question in order to show concern about the non-additive character of whatever value accrues to reliable processes in the presence of true belief. We can thus stipulate for

present purposes that token reliable processes have value whether or not they cause token true beliefs, in order to focus on the main part of the proposal.

Regarding the main point, they say,

The crucial feature of our proposed explanation is an elucidation of the psychological mechanisms whereby reliable belief-forming processes come to be accorded 'autonomous' value. Although the ascribed value of reliable processes is initially derivative from the ascribed value of the true beliefs they cause, reliable processes ultimately acquire autonomous value—value that isn't dependent, on a case-by-case basis, on the value of resultant true beliefs. (Goldman and Olsson 2009: 31)

Goldman and Olsson thus claim a special value for reliable processes, what they term 'autonomous' value.

How could this proposal help solve the Swamping Problem? I think the idea is that it is supposed to show how a reliable process can have a kind of value that is not wholly derivative on the value of true belief itself, and thus provides some hope of adding special value on top of whatever value true belief has in virtue of being true. They write,

The main possibility we suggest is that a certain type of state that *initially* has merely (type-) instrumental value eventually acquires *independent*, or *autonomous*, value status. We call such a process *value autonomization*. (Goldman and Olsson 2009: 33)

And regarding the crucial notion of autonomy at work in the proposal, they say,

Value autonomization is a psychological hypothesis, which concerns our practices of *ascribing* or *attributing* value to various states of affairs.

(Goldman and Olsson 2009: 34)

So the idea is that, though processes that produce true beliefs initially acquire all of their value from this instrumental relationship to truth, these processes come to be valued in an autonomous way or a way independent of this relationship to truth. In this way, a different kind of value is present in reliable true belief than is present in true belief alone.

Goldman and Olsson immediately notice the objection that the kind of value to be explained is objective value and that their proposal only involves an appeal to subjective valuing by individuals rather than some objective value itself. They rightly point out that it is methodologically reasonable to assume a connection between the two, on the same grounds that it is reasonable to construct theories with some (defeasible) reliance on intuitions. But this response does not address the heart of the problem. As we noted earlier, one can't read off objective

value from the mere possibility of some subjective valuing, nor can we infer the objective value of some item from the fact that some people rationally value it. Instead, something stronger is required, something like the claim that a fully informed individual not only

*might* value the item in question (in a special way) but that such an individual *would* value it in this special way.

It is crucial to note that nothing in the proposal by Goldman and Olsson provides any basis for the stronger claim. All they have is a story about how reliable processes might come to have a special kind of autonomous value. It is worth pointing out that a similar story could be told about mere statistical likelihood itself: we can imagine the process of value autonomization applying to statistically likely beliefs so that they came to have autonomous value. Doing so would force us into the position of claiming, in the original chocolate example, that the third intersection list has more objective value than the first list, the list of places that sell chocolate. Since that conclusion is mistaken, we should look for an explanation of the mistake, and the explanation is not hard to find. The explanation is that it isn't sufficient just to note that there is some psychological mechanism that could easily lead to some independent value for some property added to true belief. Instead, one needs an explanation of why a fully informed individual would autonomize in the way imagined or why one's 'better self', fully informed, would prefer that one autonomize in that way.

In this respect, the approaches of Goldman and Olsson are best thought of as in need of supplementation in order to provide any hope of success. As we will see in the next section, there is in the literature a proposal that provides just such an extension, and so we can complete our evaluation of the Goldman–Olsson proposal by turning to the appeal to final value by process reliabilists to avoid the Swamping Problem.

#### TASK 4: THE APPEAL TO FINAL VALUE

The appeal to final value arises by noting recent work in the theory of value that is ignored by those who raise the Swamping Problem. Those who raise the Swamping Problem distinguish between intrinsic value and extrinsic value, but fail to take account of the distinction between instrumental value and final value. For example, Duncan Pritchard says,

One problem with this argument is that it appears to rest on a failure to recognise an important distinction that is common currency in value theory. In particular, the value distinction that this argument ignores is that between *intrinsic value*, where an object's value is determined by its intrinsic properties, and *final value*, where we value an object for its own sake (i.e., non-instrumentally). The Swamping Problem in effect presupposes that the only way that the relational properties of a true belief—in particular, the way that the belief was formed—could contribute value would be instrumentally. But that is to

103

simply ignore the possibility that the relational properties of a true belief could contribute final value. (2007a: 6)<sup>10</sup>

Here I will put aside any concerns that there may be no such value as final value and will assume with Pritchard and others that there is a distinction needed between final value and the more familiar notions of intrinsic and extrinsic value. Since the notion of final

value is a relatively new discovery, it will be useful to have some examples of such, and Pritchard is helpful on this point:

Consider, for example, a book printed on the first ever printing press and an exact replica produced by lasers. We would undoubtedly value the former more than the latter even though we can accept, for the sake of argument, that these two objects have the same relevant intrinsic properties. Moreover, we would value the former for its own sake, quite apart from what instrumental value it had (e.g. its monetary value). The difference in value, however, is clearly due to the relational properties of the objects concerned. The reason for this difference in value is that a book printed on the first ever printing press, unlike an exact replica, has final value—it is valuable for its own sake because of how it is produced (i.e. its relational properties). (2007a: 6)

The key claim here, as far as argument goes, is the first two sentences of the first quoted passage above. The first sentence notes that usual formulations of the Swamping Problem ignore a distinction between final and intrinsic value. That point is certainly true, but by itself no more telling than noting that the argument also ignores present diplomatic tensions in the Middle East. The key argumentative move is in the second sentence: ‘In order to see that the Swamping Problem turns on a failure to recognise this distinction, one only needs to note that the problem in effect presupposes that sameness of intrinsic properties entails sameness of value.’ So ignoring the distinction between intrinsic and final value differs from ignoring certain diplomatic tensions. Why is that? Because the Swamping Problem in some way depends on a Moorean conception of value: that sameness of intrinsic properties entails sameness of value.

I want to pursue two separate lines of response to this complaint. First, I want to show that the Swamping Problem does not depend on a Moorean conception of value, and second that the distinction in question offers no hope to the process reliabilist in replying to the Swamping Problem.

So first, about the argument itself: I think it is instructive to note that as I developed the problem above, it never mentions distinctions in value at all except to distinguish purely cognitive values from other values such as those deriving from practical concerns. Given such distinctions among types of value deriving from different ends or goals, the problem as I developed it focuses on an additivity failure. Take whatever purely cognitive value is found in true belief and add statistical likelihood of truth to it, and you are guaranteed no added value.

The rest of the argument can be put in the form of a challenge to reliabilists: tell us what features beyond statistical likelihood your theory involves and how these additional features add value to the value of true belief. The first step of this argument, concerning the additivity failure for statistical likelihood surely presupposes no Moorean conception of value. Moreover, there is no reason I can see to think that the Moorean conception is involved in denying that the challenge to process reliabilists can be met. Of course, if the

Moorean thesis were true, we could explain why the non-additivity feature occurs, but we could do the same if value nihilism were true, and nobody is accusing the argument of presupposing that.

So I don't think the claim that the Swamping Problem presupposes a Moorean conception of value is correct. What is true is that some presentations of the problem have endorsed the Moorean conception, but that is not sufficient to show that an appeal to final value will somehow undermine the Swamping Problem. The Swamping Problem doesn't presuppose the Moorean view, and thus attacking the Moorean view won't allow any view to escape the Swamping Problem even if the attack on Mooreanism is successful.

Even if the Swamping Problem does not disappear on grounds of presupposing Mooreanism, however, the distinction between final and intrinsic value may yet provide the resources for escaping the problem. If we attend to the usual examples of the distinction between final and intrinsic value, things do not look promising, though. The primary example in the literature is Princess Diana's dress as opposed to a mere duplicate of it, while Pritchard's preferred example is the first book off Gutenberg's press versus a duplicate of it. Such examples take us in a direction wholly unhelpful to reliabilists. The final value involved in these cases is clearly dependent on the uniqueness of the items in question, and value depending on uniqueness of method or process is clearly a dead end in attempting to reply to the Swamping Problem.

Uniqueness, however, is not essential here, since other examples can be found where uniqueness is absent. Instead of focusing on Gutenberg's first offering or Diana's dress, we might want to focus on paintings by great artists, untouched wildernesses, first-edition books, Wedgwood china over mass-produced reproductions, and artefacts from ancient civilizations. Such examples involve rarity, but not uniqueness.

The move from uniqueness to rarity is not enough, however. While it may be true that certain pieces of knowledge are more prized than others because of their rarity, this value is not the kind threatened by the Swamping Problem. It may be true that we find wondrous and precious and fascinating the special knowledge found in theoretical physics, but this value does not generalize to our legitimate valuing of comparatively mundane knowing, such as knowing where to find a good set of tires at a decent price. The solution to the Swamping Problem has to have a kind of generality that doesn't depend on the rarity of knowledge, and perhaps we can put this point starkly by noting that it doesn't take closer and

closer approximations to the truth of radical scepticism to uncover a solution to the Meno problem. To the extent that the examples of final value depend on the rarity, there is little of comfort for reliabilists to be found in the distinction between final and intrinsic value.

But can't the reliabilist appeal here to the relative rarity of knowledge compared with the much larger class of belief itself? Such an appeal is surely correct: knowledge is rarer than belief, and rarer than true belief as well. But such a response contains the seeds of its



own destruction, for imagine a world in which knowledge is coextensive with true belief (or even belief itself). In such a world, knowledge would have no final value over and above the value of true belief or belief itself, given the assumption that final value is tied to rarity. The appeal to final value thus could not account for the special value of knowledge for such worlds, and would have to adopt the procrustean and revisionary strategy of fitting the phenomena to the theory instead of providing a theory adequate to the phenomena.

The only way to avoid this difficulty is to insist that rarity is just an artefact of the usual examples chosen, rather than an essential feature of the notion of final value. Reliabilists might note that the examples involve not only rarity, but an etiological component as well. Perhaps it is this etiological component, rather than rarity, that is the key. If such an etiological component is the key, and an etiological feature is what distinguishes justified true beliefs from unjustified true beliefs, then perhaps the final value of the former exceeds that of the latter.

This is not exactly a compelling argument: it's more like a wish and a hope. Moreover, once we think a bit more carefully about the notion of final value, we will see why the appeal to etiological features doesn't really help. We should begin by noting that final value begins with valuing in a particular way. The usual distinction between intrinsic and extrinsic value is a distinction involving the objects of value themselves, but the distinction between valuing for its own sake and valuing for the sake of something else, the distinction involved in the notion of final value, is a distinction between two types of valuing. We may want to construct some more objective notion of value out of the notion of valuing for its own sake, but we need first to notice that the fundamental notion is one concerning something we do as evaluators. The idea here is that we should focus first on different types of valuing rather than different types of values. Viewed in this way, we have two orthogonal sets of distinctions: we have things that are intrinsically valuable and things that are extrinsically valuable, and we can value such things either for their own sakes or for the sake of something else. We thus have the possibility of something of merely extrinsic value that is nonetheless valued for its own sake.<sup>11</sup>

Once we get clear in this way about the notion of final value, we can see a danger lurking along this path of appeal. In more cynical moments, we can ask ourselves what difference it really makes to the value of a thing whether Princess Diana wore it. We can understand the human desire for identification with celebrities, the associational connotations of it all, while at the same time knowing that this is all a way of coping with the lack of recognition that is part of almost everybody's mundane existence. And Guttenberg's first book is a relic, to be sure, but our interest in historical artefacts need not indicate anything really important at all, but may just be a typical over-generalization from the usefulness that attending to some of history plays in our practical lives. More generally, we know of things that have purely instrumental value that end up being valued for their own sake by distorted souls. Money is the classic example. Lovers of money of this sort

abound. They value money for its own sake, even though the only value it really has is instrumental in character. So travelling down the path of distinguishing what we value from what is valuable is worth doing, but if the kinds of valuing that would help avoid the Swamping Problem all turn out to be nothing more than all too human foibles, the appeal to final value will leave us with an unsatisfactory response to the Swamping Problem. We would have to admit that knowledge is no more important than true belief, even if we typically treat it as being so.

I do not wish to rely on such cynicism here, however, for there is a deeper worry. This deeper worry can be seen once we recognize that the fundamental notion concerning final valuing involves a way of valuing things, and then remind ourselves of what would be needed to turn this subjective notion into something objective enough to account for the special and unique value of knowledge or the special value of reliable true belief over that of true belief. For this valuing to yield any conclusion about the value of knowledge or reliability, the valuing must not be some idiosyncratic feature of some individual or group of individuals. It is, of course, possible for process reliabilists to so cherish their magnificent theory that they value the beliefs that result for their own sake on the basis of the etiology of such beliefs in processes or methods that make for objective likelihood of truth. As noted earlier, more is needed for valuing to provide the resources to explain the value of knowledge. What is needed is some way to leap the chasm from subjective valuing to objective value. We may grant to the defender of the final value of knowledge that knowledge is valued in the final way at least by some people, but how do we get from this claim to the claim that there is a value found in knowledge that is not found in, say, true belief? The point here is not the items of comparison, but rather how we get from a point that is fundamentally about a subjective state to one that is about something not subjective.

The only available answer is that some values are response-dependent values, that some things have value because they are the kind of thing that a fully

informed individual (in the right kind of circumstances) *would* value. So, the defender of the final value proposal might claim that a fully informed individual in the right circumstances would value reliable true belief in this final way, and that it is this response-dependent value that provides a solution to the Swamping Problem for reliabilism.

It is fairly easy to see how such an individual *might* value reliable true belief in this final way, but nothing in the story about reliability itself yields any basis for concluding that such an individual *would* value reliable true belief in this way. My scepticism here derives from comparing this claim with the simple statistical likelihood proposal with which we began. Adding mere statistical likelihood of truth to true belief was obviously a non-starter in explaining the value of knowledge over that of true belief itself. But there is no reason to suppose that an individual might not value statistical likelihood of truth in the final way. Such a possibility, however, is not enough. What is needed is some account

of why ordinary folk *would* value reliability in the final way, if fully informed and in the right circumstances, whereas they *would not* value statistical likelihood in that way even though fully informed and in the right circumstances. Failure to produce an account that discriminates between reliability and statistical likelihood leaves the defender of final value still susceptible to the Swamping Problem, since it was obvious from the start that the statistical likelihood proposal succumbs to the Swamping Problem.

Moreover, I doubt we have the kinds of responses needed to give any hope to this final value proposal. We simply don't have the responses predicted by the final value proposal to cases in which statistical likelihood with the specific etiological features central to reliability is added to true belief, responses that are obviously present in other cases of response-dependent value. When I look at a shard of earthenware and act uninterested, you can urge me to change my attitude by saying, 'But that's from the Mayan civilization!' My eyes light up, I'm now impressed, and I treat the shard with much greater care. When you show me a first-edition Jane Austen, I'm more reverent than otherwise. But when you tell me about a statistically likely opinion that is true, and note with appropriate fervour that its statistical likelihood is the causal result of a process that is reliable, I wonder what point you are trying to make. Why should I care about that? I just don't get it.

Of course, you can retort that there's lots of stuff I don't get. But I don't think you get it either. Think of the untouched wilderness compared to the parts of the world on which humans have already intruded, first-edition books and all the rest, Picasso paintings, and perfect replicas. When learning the final-value-imputing fact, there is an 'aha!' response (when the example is a good example of final value—nothing here turns on whether the particular contrasts I've listed really are good examples or whether other examples are more appropriate). When learning

108

that likelihood is caused by a reliable process, there is a 'ho-hum' response. It would be advantageous to the defender of the idea that there is response-dependent, anti-realist value of the sort needed to address the Swamping Problem if ordinary human beings such as we are had the 'aha!' response. But we don't, and the fact that we don't provides some evidence against the claim that the needed response-dependent value is present.

We have to be careful here, however. You may have the 'aha!' response when you compare a belief with a reliable belief, but that is the wrong comparison. Remember that the Swamping Problem is a value problem that arises once truth is in the picture. So the comparison has to be this: start with a true belief, and then add statistical likelihood of truth and then the etiological component as well. It is this comparison that, I claim, generates no systematic or uniform 'aha!' response. Once truth is in the picture here, adding reliability fails to impress.

A final point is worth making as well. The value of things is sometimes the result of features of those things themselves and sometimes the result of our caring about them.

Put somewhat differently, sometimes the direction of explanation goes from features of things to our valuing them, and sometimes the direction of explanation goes from our caring about them to the value that they have. Perhaps unspoiled wildernesses have the value they have because we care about such things, but the order of explanation goes the other way when it comes to matters involving fundamental human rights. The appeal to final value to solve the Swamping Problem presupposes that the explanation of the value of knowledge over that of true belief should be explained in one of these directions, and it is the wrong direction. I value this particular watch because it was my grandfather's, and I care about him. But I don't value knowledge like that. Knowledge is a wonderful thing, and it is because of what it is that I value it. The direction of explanation goes from knowledge and its features to my caring about it, not the other way around. Any explanation that starts the explanatory story from the way in which we value knowledge gets the order of explanation wrong.

One may resist at this point by noting the popularity of response-dependent accounts of all value, of whatever sort, according to which values always depend on what our responses are or would be.<sup>12</sup> Such accounts face the burden of avoiding the conditional fallacy,<sup>13</sup> but I will not rely on such difficulties in order

to respond to this point. For whatever the merits of response-dependent accounts of value, the point above is not a theory-driven point, but rather a datum that must be explained by any adequate theory of value. Final value is, at the initial stage, a matter of the way we value things, and the reification of this phenomenon into an item that is useful for solving the Swamping Problem involves an appeal to the ubiquity or uniformity of valuing certain things in this final way, at least in idealized conditions. The things valued in the final way fall into two classes. Some such valuing are idiosyncratic (such as treasuring one's collection of spent cartridges from some weapon) or unusual enough in some other way, so that no reification from valuing to the presence of a value is legitimate. In other cases, there is ubiquity or uniformity enough to the valuing, as with unspoiled wildernesses, relics of various sorts, and the like. There is enough uniformity in these latter cases that it is legitimate to reify, so long as we note that the value present is explanatorily dependent on the valuing in question. Even if all value is response-dependent, there is much that is valuable that does not fit this pattern, for much that is valuable is not treasured or respected enough (or is easily imagined to be so). In such a case, the nature of the explanation of value given by defenders of the response-dependence of value will be different. The explanation will need to appeal to what would be appreciated if the people in question were better informed or in different circumstances. Such an explanation is not an explanation that reifies valuing into values, but rather explains value as response-dependent without undermining the point that sometimes the order of explanation runs from caring to importance and other times it runs the other way. We might put the point this way, if we assume that a response-dependent account is acceptable. Given this assumption, sometimes the counterfactuals involved in the objective value in question are true independently of what people actually value, and

sometimes they are true because of what people value in the actual circumstances in which we find ourselves. Sometimes, that is, things would be valued in idealized circumstances because of the way they are valued in actual circumstances, and sometimes this order of explanation is unavailable or incorrect.

So, even if all value is response-dependent in some way, the datum cited remains: sometimes the explanation of value runs from features of things to our valuing them and other times from our valuing things to their value. The general difficulty raised by this datum for the final value proposal thus remains. The special and unique value of knowledge over that of true belief is a value in which the order of explanation goes from the nature and function of knowledge to our valuing it rather than the other way around. The appeal to final value explicitly reverses this order of explanation, and thus cannot give an adequate account of the value of knowledge or a solution to the Swamping Problem.

110

## CONCLUSION

The lesson to learn here is that the Swamping Problem is a much deeper problem than these solutions appear to recognize. Probabilities concerning the future don't help, an appeal to the type/token distinction is a non-starter, and appeals to final value are not successful either. The conclusion to draw is that if there is a successful defence of process reliabilism against the Swamping Problem, we have yet to see it.

111

## 6 From Epistemic Expressivism to Epistemic Inferentialism

Matthew Chrisman

### 1. INTRODUCTION

Ordinary people often make epistemic claims. For example, they claim to know the whereabouts of a misplaced book, or that someone else knows whether Jones will be at the party, and they might even do something reconstructable as claiming that Suzy's belief is justified even though false or claiming that Henry's belief is justified and true though not a case of knowledge. Insofar as they make such claims, presumably they are at least somewhat concerned to do so rightly. Epistemologists make such epistemic claims too and are, of course, also concerned to do so rightly. However, they often give priority over these epistemological issues to what I think can only be viewed as meta-epistemological issues. These are issues not about whether this or that epistemic claim is correct but about things like the standard for knowledge, the nature of epistemic justification, the possibility of rational response to sceptical arguments, and the meaning of epistemic claims.

Concerning this final issue, recent philosophical debate has largely centred on the question of whether epistemic claims are plausibly thought to be context sensitive. The default assumption has been that sentences that attribute knowledge or justification (or whatever else is epistemic) have stable truth conditions across different contexts of utterance, once any non-epistemic context sensitivity has been fixed. The idea, for instance, is that sentences of the form 'S knows that  $p$ ' have stable truth conditions once the values of S and  $p$  (plus any implicit temporal elements) are fixed. The contrary view is the contextualist thesis that such sentences do not have stable truth conditions but can vary depending on the context of utterance. So, for instance, even once we've fixed the values of S and  $p$  the sentence 'S knows that  $p$ ' could still be true when tokened in one context, while false when tokened in another context, since the truth conditions of this sentence depend in part on some

contextually varying element such as the strength of evidence required for knowledge.<sup>1</sup>

This debate manifestly presupposes that the meta-epistemological issue of accounting for the meaning of epistemic claims is to be settled by determining the truth conditions of these claims. I think this presupposition is probably false and, in any case, way too quick. This is for two interlocking reasons. First, many epistemologists see epistemic claims as evaluative or normative, in some sense.<sup>2</sup> However, in the meta-*ethical* debate—which has really become a meta-*normative* debate—most participants take alternatives to truth-conditional semantics, such as expressivism, as live options. Second, from a pragmatist view on meaning (that I think makes contact with the remit of 'social epistemology') a

natural way to approach the meaning of epistemic claims would be to ask not about their truth conditions but about what practical function they serve in our commerce with one another and other aspects of the world. It could be that, even if we agree about the truth conditions of knowledge attributions, there's still an important meta-epistemological question to be settled about their linguistic and conceptual role.

These two reasons correlate with two moments in a line of thought whose name I have taken for the title of this chapter. My purpose in it is to flesh out this line of thought and then indicate how I think it could help to enrich meta-epistemological theorizing in a way that appreciates the specially social character of epistemic claim-making.

## 2. ETHICAL EXPRESSIVISM AND EPISTEMIC EXPRESSIVISM

In the meta-ethical debate, ethical expressivists encourage us to ask not about the nature of ethical facts or values but instead about the nature of ethical evaluations. They do so because they think we can satisfactorily answer the latter question without ever countenancing some *sui generis* kind of thing—ethical facts or values—to answer the first question. Their idea is that ethical evaluations are not representations of how the world is; rather, they are expressions of our attitudes towards the way we represent the world to be. Specifically, ethical claims are seen not as descriptions of the world but as expressions of our ethical attitudes

towards things like actions, agents' characters, systems of rules, social structures, and so forth.

From the perspective of a truth-conditional semantics, which takes notions like truth and reference as its master concepts, the semantics for ethical claims that results from the standard sort of expressivism will appear quite radical.<sup>3</sup> In effect, it is the view that ethical sentences mean what they do in virtue of the state of mind that they express, rather than in virtue of the facts they purport to represent.

However, this semantical idea is not without at least some historical inspiration and contemporary allies in the philosophy of language. A rich alternative semantic programme, which is at home with this idea across the board, traces back through Grice at least to Locke. This is the 'ideationalist' or 'psychologistic' approach to meaning, which is inspired by Locke's view that 'Words in their primary or immediate Signification, stand for nothing, but the Ideas in the Mind of him that uses them' (1975: 405). In its contemporary manifestation, the project is to understand semantic content in terms of conventional regularities linking overt use of a sentence with a particular idea or state of mind.<sup>4</sup> Accordingly, all sentences are thought to mean what they do in virtue of the state of mind that they express. And, if we embrace this idea in our philosophy of language, it then opens up room for the expressivist to argue that the conventional regularities of our language are such that descriptive sentences mean what they do in virtue of expressing specific beliefs while ethical sentences mean what they do in virtue

of expressing specific desire-like moral attitudes. Because of this, I think expressivism is naturally married to an ideationalist alternative to truth-conditional semantics.<sup>5</sup>

But what exactly does the expressivist say is the desire-like moral attitude expressed by ethical sentences? The details won't matter much here, but to have a toy expressivist theory on the table, we can use a simplification of Gibbard's (1990) leading idea. That is, we can think of the sentence 'X is good' as meaning what it does in virtue of the fact that it express one's acceptance of a system of ethical norms which one believes to encourage the promotion of X.<sup>6</sup> Expressivist

114

views like this one purport to carry at least two theoretical advantages, which I'll now try to characterize in a rough and ready way.

First, expressivist views claim to have a very simple and intuitive explanation of the apparent practicality of ethical discourse. The rough idea is that the making of ethical claims seems to be intimately and causally linked with motivation to action; if, for example, one judges that giving to charity is good but displays no motivation to give to charity, then something seems to be amiss. And on a popular Humean view<sup>7</sup> motivation to action always requires the agent to have some desire-like attitude. Expressivism gives us just such an attitude in its account of what is expressed by ethical claims, which generates its explanation of the intuitive practicality of ethical claim-making.

Second, expressivist views claim to avoid queerness/placement problems for ethical facts and values. This is initially an ontological issue, but it is closely related to a conceptual issue having to do with a different sense in which ethical discourse is intuitively practical. The initial idea is that ethical facts, conceived as judgement-independent facts in the world, seem to be neither ontologically reducible to other facts whose existence is not in question nor locatable in the spatio-temporal world governed by empirically discoverable physical laws. Why? Because ethical judgements seem to have a special connection to motivation to action. Roughly, they seem to commit their author to a reason to act in a particular way, irrespective of any desire to do so. For example, if one judges that giving to charity is good, then one seems to be committed to thinking that there is a reason to give to charity. If we view such judgements as representing the world, then they purport to describe facts which, if they existed would seem to have the power to attract our wills irrespective of our desires. This can make ethical facts appear queer in the sense that they purport to have a reason-giving power not had by any other kind of fact.<sup>8</sup> Now, there is plenty of room for debate about whether this is the only plausible conception of ethical concepts and facts or whether the reduction of ethical facts to less queer facts is after all possible. However, because of the very controversies involved in this debate, any meta-ethical view ontologically committed to ethical facts or values bears a theoretical cost not borne by views that do not. Expressivism claims to avoid this cost in its contention that ethical concepts (and perhaps normative concepts more broadly) have their distinctive function in making claims that express desire-like attitudes rather than describe or represent the world.<sup>9</sup>



This is a brief outline of what ethical expressivism is and two reasons for endorsing it. It's a curious fact that very few philosophers have sought to defend or even recognized the theoretical possibility of a meta-epistemological analogue to ethical expressivism.<sup>10</sup> This is in spite of several obvious parallels. First, it's now commonplace (even if still controversial) for philosophers to insist that epistemic notions like 'knowledge' and 'justification' are normative, and many epistemologists are interested in norms of belief.<sup>11</sup> Second, some epistemologists have become quite puzzled about the nature and location of epistemic values, thinking now that if the value of knowledge were merely instrumental value for the goal of truth that it would be 'swamped' once one has the truth by whatever means.<sup>12</sup> Third, several pragmatist-inspired philosophers have asked about the social role of epistemic concepts such as 'knowledge' and come up with an account according to which this concept is not used directly to describe people but rather to do something like compliment (Rorty 1979, 1995), tag (Craig 1990), or keep the books on the ways in which various people track, process, and transmit information (Brandom 2000a: ch. 3; Rosenberg 2002: ch. 5).

In an earlier paper (Chrisman 2007), I provisionally argued for a version of epistemic expressivism not on the basis of these parallels but on the basis of the fact that it addresses two pervasive objections to epistemic contextualism while retaining the advantages of that view. I want to leave this argument to the side here, except to say that what resulted was a view that encourages us to ask not about the nature of epistemic facts and values but instead about the nature of epistemic evaluations. It did so with the typical hope of expressivist views to answer the latter question without ever needing to countenance some *sui generis* kind of thing—epistemic facts or values—to answer the former question. The idea is the same as before: epistemic evaluations are not seen as representations of how the world is; they are seen as expressions of our attitudes towards the way we represent the world to be. In particular, they involve expressions of our distinctively epistemic attitudes towards things like believers, beliefs, and types of evidence and justification.

As before, from the perspective of truth-conditional semantics, the resulting semantics for epistemic claims is apt to seem quite radical. However,

epistemic expressivists seem to have at least some support from ideationalist or psychologistic approaches to meaning on their side (as well, of course, ethical expressivists) in arguing that epistemic sentences could mean what they do in virtue of the state of mind that they express, rather than in virtue of what they represent.

The epistemic parallel to the toy version of ethical expressivism above would hold something similar about attributions of justification and knowledge. For example, it might hold that to say 'S is justified in believing that *p*' is to express one's acceptance of a system of epistemic norms which one believes to positively value S's belief that *p*. And,

sloughing over a few details, we might extend this by claiming that to say ‘S knows that *p*’ is to express one's acceptance of a system of epistemic norms, which one takes to positively value S's true belief that *p*. In addition to retaining advantages of contextualism while overcoming two of its problems (if the argument of my previous paper is correct), this view would carry at least one of the same theoretical advantages as ethical expressivism.

This is the second advantage mentioned above: Epistemic expressivism would avoid queerness/placement problems for epistemic facts and concepts. Philosophers have been less concerned about the ontological status of epistemic facts and values than ethical facts and values; however, arguably, whatever naturalistic inclinations motivate concern about the queerness/placement of ethical facts or concepts, the same naturalistic inclinations should motivate the same concerns about epistemic facts and concepts, insofar as we're conceiving of these as genuinely normative (whatever, exactly *that* means). If that is right, then epistemic expressivism would avoid an ontological/conceptual cost in its contention that epistemic concepts function like ethical concepts to make claims that express attitudes rather than describe or represent the world.

The other advantage typically claimed for ethical expressivism was its intuitive explanation of the practicality of ethical discourse, and it is perhaps because epistemic discourse isn't seen as similarly practical that epistemologists have ignored that accounts of the meaning of epistemic claims that don't turn on the truth conditions of these claims, such as expressivist accounts. After all, claiming, for example, that Plato knew nothing about Google does not seem to be as tightly connected to being motivated to act in a particular way as claiming, for example, that charity is good. However, although I think it's probably correct that epistemic discourse is not practical in exactly the same way that expressivists claim that ethical discourse is practical, I suspect that epistemic discourse is interestingly practical in another sense that encourages rethinking the pervasive presupposition in epistemology about how best to approach an account of the meaning of epistemic claims.

More on this in §4, but first I want to discuss another—to my mind, better and more general—reason for being wary of epistemic expressivism as an account

117

of the meaning of epistemic claims.<sup>13</sup> A brief and somewhat dogmatic discussion of the notorious Frege–Geach problem in meta-ethics will help to expose this reason.

### 3. EXPRESSIVISM'S DISCONTENTS

There are other views going under the label ‘expressivism’ on the meta-normative scene<sup>14</sup>, but one of the main challenges for views, which hold that normative claims mean what they do in virtue of the type of desire-like attitude that they express, is to explain, in a systematic way, the logic of normative concepts in all of their uses. Normative concepts are clearly used not only atomically but also in logically complex contexts. We say things not only like ‘X is good’ but also like ‘If X is good, then Y is good’; not only like ‘S's

belief that  $p$  is justified' but also like 'If S's belief that  $p$  is justified, then S's belief that  $q$  is not justified'. If expressivists say that the atomic claims mean what they do in virtue of expressing a desire-like attitude, what will they say about the logically complex claims? It seems that one could endorse these claims without expressing a desire-like attitude (towards X or Y, or towards S's belief that  $p$  or S's belief that  $q$ ). This is the Frege–Geach problem (Geach 1965).

In my view, ethical expressivism in its early forms suffered the Frege–Geach problem because, in effect, it treated the meaning of normative predicates as given completely by the expressive force of applying them in making an atomic normative claim. This is why Ayer suggested that atomic moral claims like 'Stealing money is wrong' have no factual meaning and express no proposition. But such a view is hopeless when it comes to explaining the semantic properties of normative predicates embedded in force-stripping contexts such as conditionals. There must be some stable content between 'Stealing money is wrong' and the antecedent of 'If stealing money is wrong, then it will be good thieves are punished', if we're to view a modus ponens inference proceeding from these premises to be innocent of the fallacy of equivocation.

Interestingly, Ayer's view of complex normative sentences of another type may, if suitably generalized, provide some traction against this problem. For the case of a complex normative sentence, he suggested that the sentence has descriptive content that is presented in a particular way. His example is the claim 'Your stealing that money was wrong', and he suggests that it expresses the proposition that you stole that money, but it presents this proposition in a special way that

118

could in principle be replaced by a special form of exclamation marks indicating a speaker's contempt for the action mentioned. What is interesting is that this can help to capture some stable content across unembedded and embedded contexts for such complex claims. For Ayer can say that the same proposition is the content of the antecedent of 'If your stealing that money was wrong, then you will be punished'.

A problem with this, of course, is that we need to explain the content of not only complex moral claims but also atomic moral claims. But many latter-day expressivists have sought to overcome this problem by also holding that atomic claims have *both* descriptive and evaluative aspects.<sup>15</sup> It is highly controversial whether any of these enriched expressivist views work, but, rather than get into the trenches of that debate, I want to point out that, insofar as these views continue to operate within the general ideationalist approach to meaning, they can still be fit into Ayer's heuristic at least this far: normative sentences are still viewed as having a particular descriptive content that is presented in a special normative way that could in principle be replaced with a special set of exclamation marks expressing the relevant desire-like attitude.

For example, we can already see how this is the case with our toy version of ethical expressivism above. The idea was that a sentence of the form

‘X is good’

expresses the acceptance of a set of norms N, which one takes to encourage promoting X. Here the descriptive content is that

N encourages promoting X

and it is presented, so to speak, in the mode of one who accepts N. Thus, on the view, it seems that we could in principle replace the claim ‘X is good’ with ‘N encourages promoting X!!’, where the ‘!!’ indicates that the speaker accepts

119

N. When the normative predicate is embedded in a force-stripping context, the point remains the same. For instance, the sentence

‘If X is good, then Y is good’

might be thought to have the descriptive content that

if N encourages promoting X, then N encourages promoting Y

which is presented in the mode of one who accepts N. Thus, it could, in principle be replaced with ‘If N encourages promoting X, then N encourages promoting Y!!’, where the ‘!!’ indicates that the speaker accepts N.

No mature version of expressivism deploys Ayer's heuristic tactic so baldly, but I think most can be fit into the general model.<sup>16</sup> Because of this, however, I worry that they trade on an impoverished view of what a normative concept could be. To appreciate this, let me contrast two conceptions of concepts.<sup>17</sup>

There is a strong tradition in cognitive science of viewing concepts, especially those often deployed in observation, as asymmetrically dependent labels or classificatory devices. On this view, concept possession is a discriminatory ability: to possess the concept C is to be reliably disposed to respond differentially to things in a particular class, the C-things. For instance, the concept ‘red’ is, on this view, a classificatory device for the red things; to possess this concept is to be reliably disposed to respond differentially to red things (under normal circumstances and with the right sort of prompting) with the word ‘red’ (or some suitable analogue), that is, to be able to classify things as red or non-red.<sup>18</sup>

Another tradition dominates the philosophy of logic, where concepts are seen as logically articulable functions or nodes of inferential (evidential and/or justificatory) relations. On this view, concept possession is a rational ability: to possess the concept C is to be able to draw the right sorts of inferences to and from applications of the concept C. For instance, the concept of conjunction is, on this view, the node of the inferential relations between judgements deploying this and other concepts; to possess this concept is to have the ability to make

120

(some substantial portion) of these inferences—for example, from ‘P and Q’ to ‘Q’ and from ‘P’ and ‘Q’ to ‘P and Q’.<sup>19</sup>

Which conception of concepts is correct? I think both conceptions capture something that is important about at least some concepts, so I'd say neither is correct if it is meant to capture all that there is to concepts and concept possession. However, when it comes to the nature of *human* concepts and especially our normative concepts, I think the inferentialist conception has some claim to being more illuminating and fundamental. For one thing, it doesn't have to *deny* that concepts sometimes serve as classificatory devices, it just claims that there is more to human concepts. For another, when concepts have a richer logic than is captured by their possible use as classificatory devices, this view can explain this logic. This seems to be especially apt when it comes to logical concepts such as ‘necessary’, but it may extend to other concepts such as normative concepts too. Even those concepts that are unquestionably observational might be thought to embody inferential nodes. For example, it's plausible to think ‘X is red’ implies ‘X is colored’ and ‘X is crimson’ implies ‘X is red’. And this could be relevant for determining their classificatory function.

I mention this distinction between the two views of concepts because I suspect that expressivists are, to put things somewhat crudely, tacitly conceiving of normative concepts as classificatory devices rather than as nodes of inferential relations. It's not that they think that a normative concept such as ‘good’ is analogous to an observational concept such as ‘red’ in bearing an asymmetric relationship to all and only the good things. That would be a form of realist cognitivism. Rather, it seems to me that they treat ethical concepts such as ‘good’ as partially second-order devices for classifying the things towards which we (the concept-users) have a certain desire-like attitude in virtue of certain descriptive properties of the things.

In other words, although, on the expressivist view, ‘X is good’ clearly does not describe X as something the speaker approves of, it does classify X as one of the things to which the speaker has a specific desire-like attitude of approval. The same is true even in logically complex contexts such as a conditional. For example, on the expressivist view of the sentence ‘If X is good, then Y is good’, neither X nor Y is classified as unconditionally good. But they are classified as standing in a certain relation with respect to the things the speaker approves of—basically, if X is in the class of things the speaker approves of then so is Y. This is why we can reconstruct the expressivist's view of normative sentences as one according to which they have descriptive content that is presented in a special way, which we could in principle replace with a certain kind of tag or label, for example, the special exclamation marks from Ayer.

On this way of thinking of things, to possess a normative concept is to be disposed to respond differentially to the things towards which one has a certain desire-like attitude (under normal circumstances and with the right sort of prompting) with the word ‘good’ (or some suitable analogue). This is why the expressivist has such an easy time

explaining the intuitive practicality of normative discourse, as explained above. However, just as I think the general conception of concepts as classificatory device will have a hard time explaining, in general, the logic of human concepts, I think the view of ethical concepts which seems to be tacit in expressivism, is likewise bound to have a hard time explaining the logic of ethical concepts.

There may be a structurally adequate expressivist response to the Frege–Geach problem. That is to say that there may be a version of the expressivist idea that normative claims express desire-like attitudes, which is capable of assigning different mental states to different sentences as a way of giving their meaning at whatever degree of logical complexity. However, as long as the aspect of normative sentences that makes them normative is conceived as a second-order desire-like classificatory device, I think expressivists will be stuck with an impoverished view of normative concepts. This is not to say that it is wrong to view normative concepts as (partially) second-order classificatory devices, for that idea may very well capture an important aspect of these concepts. However, especially in the normative case, it seems to me that this cannot be the most illuminating or fundamental story, since it does not explain the inferential (evidential or justificatory) relations in which normative concepts stand.

This is, I think, a better reason for being wary of expressivism, and it may seem to be even more pressing in the case of epistemic expressivism than in the case of ethical expressivism. For perhaps philosophy is still influenced by positivist undercurrents enough to take seriously the idea that ethical concepts are somehow attitudinal classificatory devices rather than inferential nodes. However, positivism's influence hasn't extended far enough to make the idea that epistemic concepts are attitudinal classificatory devices rather than inferential nodes equally acceptable. Epistemic concepts clearly have a robust logic, and it seems to be one entirely parallel to other non-normative concepts. This is something that I think any view of normative concepts, which turns implicitly on the cognitive science conception of concepts rather than the philosophy of logic conception of concepts, is bound to fail to explain.

#### 4. INFERENCEALISM AS A NEW META-ETHICAL AND META-EPISTEMOLOGICAL VIEW

Summing up before moving on: meta-epistemology has largely presupposed that debates about the meaning of epistemic claims will be settled by determining the truth conditions of these claims. I have claimed that this obscures meta-normative

views like expressivism which explain the meaning of the target claims in terms other than truth conditions. However, in both meta-ethics and meta-epistemology, although expressivism claims certain theoretical advantages, it also seems to tacitly assume a view of normative concepts that will not be taken to be the most fundamental story in the epistemic case and, anyway, is impoverished enough to wonder whether it should be taken as the most fundamental account of any concepts at all.

In the light of the distinction between the two views of concepts, I now want to suggest a way to retain the advantages of expressivism without incurring the problem I've raised for it. The strategy will be to pick up on the idea of concepts as inferential nodes and to sketch the rudiments of a new meta-normative view with this at its heart. As a first pass, this view can be seen as involving three related elements.

First, it says that we should think of normative concepts not (merely) as classificatory devices (even at the second-order) but as nodes of inferential relations. Which inferential relations? It is difficult to say in general and exhaustively, but these will comprise all of the standard logical relations (like *modus ponens* and *modus tollens* inferences) and the conceptual connections between these concepts and other concepts (e.g. ought implies can, knowledge implies truth, etc.).<sup>20</sup>

Second, it says that we should recognize the distinctively practical kinds of inferential relations in which some normative concepts are caught up. For example, while 'S ought all things considered to do A' may have the theoretical implication that 'S can do A', from the first-person, 'I ought all things considered to do A' may have the practical implication 'I shall do A' (where this is not the expression of a prediction but rather the expression of an intention). Likewise, while 'S knows that *p*' may have the theoretical implication 'S believes that *p*'; from the first person, 'I know that *p*' may have the practical implication 'I shall act as if *p* is true' and/or 'I shall stop enquiring as to whether *p*' (where this is not the expression of a prediction but rather the expression of an intention). In effect, here we get a distinction between two different sorts of inferential roles: the theoretical inferential roles that normative concepts may share with descriptive concepts and the practical inferential roles that mark them out as different.

Third, it says that we should retain, if only provisionally, the popular idea that intention and belief have different directions of fit with the world, such that, if two people believe incompatible propositions to be true, their disagreement presupposes some fact of the matter about who is right, whereas, if two people intend to make different propositions true, their disagreement does not by itself presuppose some fact of the matter about who is right.

Now, if we hold that normative predicates and the sentences in which they figure mean what they do in virtue of their inferential role, so construed, I think we get an alternative to both truth-conditionalism and expressivism, as accounts of the meaning of normative claims. We might call this view *normative inferentialism* (with *ethical inferentialism* and *epistemic inferentialism* as two species).<sup>21</sup> Its attractiveness depends on its ability to retain the advantages of expressivism while enriching the conception of normative concepts in a way that avoids expressivism's problems. I've just said how it enriches the conception of normative concepts by moving beyond classificatory devices to inferential nodes. I'll illustrate separately how it might retain the advantages of ethical expressivism and then turn to epistemic expressivism.

In the meta-ethical debate, expressivists claim an easy explanation of the intuitive practicality of ethical discourse. However, above, I distinguished two senses in which ethical discourse is distinctively practical. There is an issue in, so to speak, the order of causes: those who make certain sorts of ethical claims tend to exhibit correlated motivations. By contrast, there is an issue in the order of reasons: one can, it seems, make ethical claims to justify actions. I think ethical inferentialism is consistent with many different explanations of the first sense in which ethical discourse is practical. Indeed, it is consistent with the idea that ethical claims express desire-like attitudes that serve, in the order of causes, to motivate the expected action (although, it will not put this idea to the same work in the theory of meaning that the expressivist does). However, ethical inferentialism offers a novel explanation of the second sense in which ethical discourse is distinctively practical. By treating ethical predicates as standing in inferential relations with the predicates that serve in the expression of intentions, we can explain more than just why ethical claim-making tends to cause certain motivation, we can explain why it can be used in an attempt to justify it. The idea is that it is by standing in inferential relations to intentions (to action) that ethical claims are of the right inferential category to be made in response to demands for justification of an action. For example, if you give some money to charity, and then someone asks you, ‘Why did you do that?’, you can typically say, ‘Because giving to charity is a good thing to do’. That is, you can answer, in the order of reasons, by making an ethical claim. Ethical inferentialism explains why this is so.<sup>22</sup>

The other main advantage of ethical expressivism was its claim to avoid queerness/placement worries generated by the practicality of ethical concepts.

124

Here, it's difficult to assess the relative merit of inferentialism because it's difficult to say when the inferentialist semantic idea, viz. that claims mean what they do in virtue of their inferential roles, commits one ontologically. Some inferentialists (e.g. Dummett) seem attracted to anti-realism across the board, in which case ethical inferentialism would obviously be just as ontologically parsimonious as ethical expressivism, but only on pain of deconstructing the whole issue of ontology. Other inferentialists (e.g. Sellars, Harman) seem willing to reconstitute the notion of representing the world within the inferentialist framework. However exactly they do this, as long as there is room for viewing some claims as not completely representational because of their inferential relations to expressions of intentions, ethical inferentialism can argue that ethical claims fit into this latter category and thereby retain the advantage of ontological parsimony claimed by ethical expressivists. The idea would be to think of some concepts as bearing no direct inferential connection to intentions, and so not having a distinctively practical inferential role. These concepts could be viewed as the descriptive or representational concepts, whose deployment in basic assertions commits the speaker to the existence of the corresponding fact.<sup>23</sup> For example, ‘X is red’ could then be viewed as committing one to the fact that X is red. By contrast, other concepts could then be thought of as lacking this commitment precisely because of their distinctively practical inferential role (and the



corresponding difference in ‘directions of fit’). Assuming that ethical concepts belong in the latter category, ethical inferentialism would have a parallel way to avoid queerness/placement worries to the ethical expressivist.<sup>24</sup>

The ostensible advantages of an expressivist approach have not seemed as weighty in the meta-epistemological debate about the meaning of epistemic claims as they may seem in the meta-ethical debate. However, in the case of queerness/placement worries, I think this may have resulted from mere oversight. If we conceive of them as normative, epistemic facts and values seem to me to be just as difficult to locate in the natural world as ethical facts and values. For one thing, epistemic concepts seem to resist analysis in non-normative terms. For another, epistemic norms seem to be categorical and not hypothetical, in the sense that their normative force does not depend on the contingent and personal aims or desires of the agents to whom they apply.<sup>25</sup> In any case, to the extent that *ethical* inferentialism retains resources for claiming ontological parsimony over truth-conditionalist views of the meaning of ethical claims, *epistemic* inferentialism

125

will do so as well. So, if epistemic facts and values are queer enough to make ontological commitment to them a theoretical liability, epistemic inferentialism will retain some of the advantage of epistemic expressivism.

As I mentioned above, the expressivist's easy explanation of the motivational force of ethical discourse, does not seem to carry over to epistemic discourse. We can, of course, make some pretty reliable predictions about what someone who affirms certain epistemic claims will do. For example, there may be a tight connection between claiming ‘I know that *p*’ and being motivated to act as if *p* is true. However, it seems that this has nothing specifically to do with the claim's being an epistemic claim. We can just as well make such predictions from someone who asserts that *p*.<sup>26</sup>

Nonetheless, it seems to me that epistemic discourse is practical in a different sense. As with ethical claims, I think epistemic claims can be made in the attempt to justify certain actions. Which epistemic claims can be used to justify which actions is a matter of considerable debate. Some epistemologists think knowing that *p* justifies asserting that *p* or acting as if *p* is true, while others think being justified in believing *p* is enough to justify these actions.<sup>27</sup> Whatever exactly the correct account is, if some epistemic notions serve as the norm of some kinds of actions, then epistemic discourse exhibits a sort of practicality in the order of reasons that demands explanation. Epistemic inferentialism provides this explanation by construing epistemic predicates as inferentially related to the predicates by which we express intentions.

Thus, it seems to me that normative inferentialism—with its ethical and epistemic species—will be at least as well off as normative expressivism vis-à-vis the alleged advantages of the latter view. And, since normative inferentialism is based on a view of concepts, including especially ethical and epistemic concepts, as inferential nodes rather

than classificatory devices, it will overcome one of the core drawbacks of normative expressivism: its impoverished conception of normative concepts.

## 5. TOWARDS A VIABLE SOCIAL EPISTEMOLOGY IN PRAGMATISM

In many ways, the sort of epistemic inferentialism for which I have been arguing here is inspired by a sort of pragmatism about words and concepts that encourages us to look for their social-practical roles rather than referents in

126

order to understand their meaning. However, it's a pragmatism which makes a distinction between normative concepts like 'knows' or 'good' and descriptive concepts like 'water' or 'heat'. The difference will be not in which concepts have inferential connections but in the inferential connections between these concepts and action. In my view, normative concepts embed the concept 'ought' and this we should understand in terms of its inferential connections to intentions to action. I think this allows us to view normative concepts as playing some essentially non-descriptive social-practical role.

Because other pragmatists sometimes say things sounding roughly like this, I'd like to conclude by distinguishing the pragmatist social epistemology an epistemic inferentialist should allay himself with from some other similar ideas.

As I mentioned above, pragmatists often challenge the standard approach to doing meta-epistemology. For example, Rorty (1979, 1995) suggests that epistemic concepts are primarily terms of praise. This allows him to cast the endless debate about the nature of knowledge as stemming from a failure to recognize that 'knows', like 'good', is not a word that picks out something with a nature. In a somewhat different vein, Craig (1990) proposes to approach the concept of knowledge by investigating what role the introduction of such a concept would have served in human coordination in community, much like Hobbes tried to understand the concept of justice by investigating what role the introduction of such a concept would have served in human coordination in community. Thus, Craig's account of 'knows' construes it as (initially) a way to tag reliable informants; what it takes for someone to be deserving of such a tag will largely be down to the varying specific information tracking and transmitting needs of different groups of people and not some underlying 'nature of knowledge'.<sup>28</sup>

Both of these approaches to understanding epistemic concepts have met significant resistance, which is I think due to the vague sense that they undermine the objectivity of epistemic concepts and discourse. If epistemic terms are used merely to praise people's beliefs in a special epistemic way, and different people find different sorts of things to be worthy of epistemic praise, then there will be no objective answer to first-order epistemic questions. Similarly, if epistemic terms are used to tag reliable informants, but there is no specific measure of reliability one must achieve to deserve such a tag, then there will be no objective standard for our epistemic claims.

Despite this worry, I think it would be wrong to doubt that epistemic concepts are used to praise and tag; it's just that they seem to do or be something more. If the view of

epistemic concepts as devices of praise or tagging encourages something like the expressivist semantics for epistemic claims that I discussed

127

earlier, I think the inferentialist semantics that I have proposed to view as expressivism's successor (both in meta-ethics and meta-epistemology) offers a more nuanced view of the essentially social function of epistemic concepts and, thereby, a way to view them as doing more than merely praising and tagging.

The idea would be to treat epistemic concepts as nodes of inferential relations, which are ossified and perpetuated by the social-linguistic practice of users of these concepts in human linguistic community as they faced certain needs. Thus, there will be good pragmatic reasons to explain why it is better for us to use concepts embodying these inferential roles rather than some other concepts embodying some other inferential roles. However, once we are using the concepts we have good reason to use, since they embody specific inferential roles, there is an objective answer to questions about when they can be correctly applied and what follows from such an application. (Compare: an inferentialist about logical concepts usually holds that there is an objective answer about the correct application of, and what follows from, a conjunction or negation.) This is why I have taken my inspiration from other pragmatists about knowledge, like Brandom (2000a: ch. 3) and Rosenberg (2002: chs. 5–6) who embed their pragmatism about epistemic concepts in a broader pragmatism about concept possession and use more generally that underwrites an inferentialist treatment of these concepts.

The objectivity for which this move makes room is not the objectivity of some sort of realism about the relevant kinds of predicates and concepts that they express (e.g. moral, epistemic, logical, etc.). Rather, it's the kind of objectivity sought by Kant's discussion of pretty much all of the concepts that interested him. His idea is that we can establish our unconditional right to use the concepts without needing to establish access to properties in the world that these concepts purport to represent. The value I see in treating this as stemming from the social function of epistemic discourse is that it encourages investigation into why it is useful to use epistemic concepts the way that we do. And this is the sort of investigation that single-minded focus on the truth conditions of epistemic claims (invariantist or contextualist) or even on the expressive role of epistemic claims (representationalist or expressivist) is likely to hinder rather than help.<sup>29</sup>

128

## 7 Norms of Trust

Paul Faulkner

1.

Norms are instructions how to act. Their form is the imperative or hypothetical imperative: ‘Do X’, ‘Don’t do X’, or ‘if Y, then do X’. There are moral norms, epistemic norms, norms of practical rationality, and social norms. In following norms we come to believe truths, to be justified in our belief, act rationally, lead a virtuous life, and act in morally permissible ways. In following social norms we come to lead a life that is acceptable to a given society, a life that conforms to this socially established way of living. Social norms thereby differ from other norms in that their prescriptions are relative rather than universal. At the most trivial level these norms can be no more than a matter of etiquette: they are a matter of adopting the right register, wearing the right clothes, or, for instance, using the outermost knife and fork first. In these matters most would accept that when in Rome one should do as the Romans do. But this easy-going relativism towards the prescriptions of social norms is not always so easy: social norms can be strongly felt. And nor are all social norms trivial: adopting the right register, for instance, can shade into treating someone the right way and thereby engage moral norms. Similarly, prescriptions of politeness can overlap with epistemic prescriptions. ‘Believe people’ might be a matter of politeness, but it could also be an epistemic norm if truth-telling were the norm. Is ‘tell the truth’ a social norm, and so the norm? Should we believe people? This chapter aims to argue that, roughly speaking, we have these social norms. It aims to offer an explanation of these social norms and an account of how this bears on epistemological theories of testimony as a source of knowledge.

2.

The attitudinal hallmark of social norms, Jon Elster suggests, is that they are associated with, and sustained by, ‘feelings of embarrassment, anxiety, guilt and shame that a person suffers at the prospect of violating them, or at least at

129

the prospect of being caught violating them. Social norms have a *grip on the mind* that is due to the strong emotions their violations can trigger’.<sup>1</sup> Where the social norm concerns the behaviour of one party towards another, there are three attitudinal dimensions that could be individuated. Violation of the norm will provoke emotions of shame or guilt in the ‘wrongdoer’; it will provoke the reactive attitude of resentment in the ‘wronged’; and it will provoke punitive attitudes of disapproval or anger in third parties. These hallmark emotions are found in our attitudes towards truth-telling and believing others. This might be illustrated for truth-telling as follows. A native to this city you are approached by someone who is visibly a tourist and asked directions to the train station. Most would think it would be the wrong thing to do to misdirect the tourist, and that it would be shameful to misdirect the tourist ‘for the fun of it’. A friend X has confided in you. The

matter is of some delicacy. Another friend Y is curious about what is going on with X, you try to avoid the issue but Y is persistent and asks about the matter directly. With no room to manoeuvre you choose to maintain X's confidence and lie to Y all the while chaffing at being put in this position. These cases illustrate, at the very least, that in the context of being quizzed for information we feel that we *should* give the audience the information needed. Equally, we would be susceptible to guilt-like emotions were we to mislead the audience—even if for good reason as in the case of the nosey friend. And we feel it would be appropriate for the misled audience to resent being misled. These hallmark emotions then suggest that we have some kind of norm of truth-telling. That there is a parallel norm of believing others is then illustrated by imagining things from the other side. Suppose that you look like the kind of person whose directions would be authoritative but after having given clear and confident directions to the train station you witness the tourist walk off in the opposing direction and promptly ask someone else for the same set of directions. This manifest distrust has something of an insult to it.<sup>2</sup> If humour did not intervene, it would be liable to provoke something like resentment: what reason could he have had for not believing you? And one would expect the tourist to be embarrassed if he sees you watching him. This set of emotional responses equally suggests that we have something like a norm of believing others. So the first question is what is the content of these two norms?

The 'norm of truth-telling', I think, is less one of truth-telling, and more one of being cooperative in conversation. Suppose that as an audience you need to know whether *p* and engage a speaker in conversation with the purpose of finding this out. Ideally, you want to engage with a speaker who will tell you that *p* if *p* and tell you that not-*p* if not-*p*. Ideally you want a speaker who makes her conversational contribution one that is true. But since it is not always plain

130

what is true, the most one can really expect is for a speaker not to say what she believes is false or for which she lacks adequate evidence. But this is not all. Williams gives the example of being told 'Someone's been opening your mail', when it is the speaker who has been doing so. This speaker has said something true, but what she has said is misleading because it implies the falsehood that *someone else* has been opening the mail. To communicate a truth, she would have to say more: 'Someone has been opening your mail and that someone is me' (or 'I've been opening your mail'). So if you want to know whether *p*, then, in addition to wanting a speaker to try and say what is true, you also want the speaker to be as informative as is required for your not being misled. Ideally the speaker's contribution would also be appropriately relevant and lucid. And if it is all of these things, the speaker's contribution has been guided by Paul Grice's maxims of Quality, Quantity, Relation, and Manner respectively. Given our conversational goal of learning whether *p*, a reply that is guided by these maxims is cooperative.

On the other side, the norm of believing others is less one of credulity and more the paired norm of presuming cooperation.<sup>3</sup> Grice's claim is that we can presume

conversation to be cooperative; we can expect it to be guided by 'the Cooperative Principle':

We might then formulate a rough general principle which participants will be expected (*ceteris paribus*) to observe, namely: make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. One might label this the Cooperative Principle. (Grice 1967: 26)

For Grice, it is the presumption that participants are following this principle that makes a 'talk exchange' a 'conversation' rather than a 'succession of disconnected remarks'. This presumption, his point was to observe, can be needed to work out what a speaker means by what she says. For example, a speaker *S* states 'You're a fine friend' on learning that her close friend *A* has divulged her secret to a business rival. In knowing that *S* knows of his actions, *A* knows that *S* believes that what she says is false. So *S* appears to be flouting the maxim of Quality and with it the Cooperative Principle. But they still seem to be having a conversation: *S* appears to be telling him something. So *A* must presume that *S* is following the Cooperative Principle and is telling him the opposite of what she says. There is little cooperation between *A* and *S* but *A*'s receiving the information from *S*'s telling that he is a poor friend requires the conversation be a cooperative endeavour. Now, *not all* talk exchanges are cooperative: not all talk exchanges are *conversations* in Grice's sense. We might expect our interlocutors to cooperate in conversation, but they need not. This is illustrated by the mail case. The speaker

131

*S* purportedly tells the audience *A* something. Since the conversation thereby has the seeming purpose of *S* giving *A* some information, *A* will expect *S*'s utterance to be such as is required, and to be true and appropriately informative. On the presumption that *S* is following the cooperative principle, *A* will thereby understand *S* to be telling him that someone else has been opening his mail. And there is room here for *S* to say truthfully that she didn't say *that*, and so has been misunderstood. She has been misunderstood because the presumption of cooperation is false; it turns out that they are not having a conversation in Grice's sense: *A* thinks that he is being told something when in fact he is being manipulated. But *A* was right to presume that they were having a conversation—*A* was right to presume the Cooperative Principle was being followed—not merely because this is how things seemed, but also because the Cooperative Principle is a normative principle amounting to the prescription that if you want to have a conversation which has a certain understood purpose, then you'd better make your conversational contribution such as is required by this accepted purpose. If the accepted purpose is the giving and receiving of information, then you'd better try and say what is true and be appropriately informative.

When a speaker satisfies these maxims, let me say that the speaker is *trustworthy*. Grice's proposal that participants in a talk exchange should follow and be presumed to follow the Cooperative Principle is then the proposal that we expect speakers to be trustworthy

when the talk exchange has the accepted purpose of giving and receiving information. The idea that conversation be seen as a cooperative endeavour thereby yields a pair of social norms. The prescription that speakers follow the Cooperative Principle and its maxims describes a *social norm of trustworthiness*. And the paired prescription that as audiences we presume this of speakers and act as if we believe that they are following the Cooperative Principle and its maxims describes a *social norm of trust*. Together this pair of norms describes a standard that we expect interlocutors to live up to when engaged in a certain practice: that of having a certain type of conversation. On this standard if another depends on you for information, then other things being equal you should try to say what is true and try to be appropriately informative; and if another purports to tell you that something is so, then other things being equal you should explain this in terms of their trying to be appropriately informative. The fact that the presumption that speakers follow the Cooperative Principle reveals what speakers mean by what they say so in a way that naturally describes how we understand others I then take to be good evidence that this pair of norms describes our conversational practices. Grice's definition of a conversation is not meant to be a mere term of art. This answers the first question of the content of the norms of truth-telling and belief: they are actually norms of trustworthiness and trust. The next question is what explains our having these social norms? What explains the fact that talk exchanges tend to have the civility of conversations?

132

3.

An influential account of social norms is provided by David Lewis (1969). Lewis's starting point is the idea of a *coordination problem*. The resolution of this problem, Lewis argues, gives rise to conventions which then define social norms. We often need to coordinate our actions. A simple case: if we live in the suburbs on opposite sides of the city and decided to meet for a drink after work in the city we need to arrange a time and a place. Supposing neither of us cares whether we meet at six or seven or at The Lion or The Lamb, then we have what Lewis defines as a coordination problem: we have a coincidence of interests in that of the four possible time and pub combinations it doesn't matter to either of us which pair is chosen but each likes one pair best given the other's choice. The problem is settling on a choice. If you choose seven at The Lamb that is fine by me and best for me, and we could reach this arrangement by declaration and agreement. However, another possibility is coordination by precedence. If you get cut off just as you were about to make this suggestion, one arrangement may remain salient and this is a repeat of last week's meeting. If this arrangement then successfully repeats—we both turn up at the same time and place—it will establish an expectation of future conformity and we will have an embryonic convention: to meet at The Lamb at seven after work on Wednesdays. This gives Lewis's first approximation:

A regularity R in the behaviour of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- (1) everyone conforms to R;
- (2) everyone expects everyone else to conform to R;
- (3) everyone prefers to conform to R on condition that the others do, since S is a coordination problem and conformity to R is a proper coordination equilibrium in S. (Lewis 1969: 42)

This definition does not contain any normative terms, however social norms, Lewis suggests, can be defined as ‘regularities to which we believe one ought to conform’ (1969: 97), and conventions are norms in this sense. This ‘ought’ Lewis explains in two ways. First, conformity to a convention is in everyone's interest: the regularity is a proper coordination equilibrium or a combination that each likes better than any other, given the others' choices. And, on Lewis's refinement, it is *common knowledge* that this is so (1969: 52 ff). Consequently, everyone will recognize two reasons why one ought to conform: one ought to do what is in one's interests; and one ought to do what others expect one to do when this expectation is reasonable. And this expectation is reasonable given that conformity is to a regularity that everyone knows is in everyone's best interest.

133

The ‘ought’ of social norms then carries the force of these two reasons on Lewis's account. It also has further force, which is detailed by Lewis's second explanation: the ‘ought’ of social norms is attached to feelings of approval and disapproval. If others are confronted by an action that fits with their reasonable expectations, then they are likely to approve; and if others confront an action that runs counter to these expectations, then they are likely to explain it discredibly. So a failure to conform to a convention elicits disapproval, and this constitutes a sanction. This makes conventions ‘by definition, a socially enforced norm: one is expected to conform, and failure to conform tends to provoke unfavorable responses from others’ (Lewis 1969: 99). Sanctions offer additional motivation for compliance. So conventions are social norms because they have normative force: one has reasons to comply with them *and* there are sanctions on one's compliance. This explains why the regularities that constitute the norm persists. What explains there being regularity in the first place is that norms as conventions are proper coordination equilibria, and so in everyone's best interest.

4.

In order to give this Lewisian explanation of the Cooperative Principle, and more specifically the norms of trust and trustworthiness, these norms must prescribe courses of action in a situation whose outcome is determined jointly by the actions of two or more. This they do. This situation—call it the *testimonial situation*—is a conversation whose ostensible purpose is the giving and receiving of information. The prescribed outcome is for the speaker to be trustworthy and so to try to tell the truth in an informative way and for the audience to trust and so to act as if he believed the speaker is doing just this. In order for these norms to have the status of conventions, this outcome must be a proper coordination equilibria: it must be an outcome that each likes best given the other's



choice. Now, in any particular case this could be true. But conventions and norms must hold independently of the particularities of any given case. And it is easy to imagine cases where this outcome is not a proper coordination equilibria. Holding the audience's trust constant, a speaker might have preferred the outcome where he was untrustworthy. Suppose the speaker has just made a kill in a hunt and would rather keep this kill for himself and his family. In this case, if another hunter asks him whether he had any luck hunting, he might well tell the other hunter the truth, but he would prefer the outcome where there is trust but he keeps the fact that he has made a kill concealed. So in this particular case the prescribed outcome is not a coordination equilibria. Moreover, if the testimonial situation is considered in abstract there is no reason to think that this case is peculiar. The recurrent situation is that of a conversation whose ostensible purpose is the giving and receiving of information, wherein the audience ostensibly needs to know whether  $p$ , and the speaker ostensibly tells the

134

audience what he needs to know. The choices for the audience are to trust or not—to accept what the speaker tells him on the presumption that the speaker is trustworthy or not; and for the speaker to be trustworthy or not—to try to tell the audience the truth informatively or not. Given the assumptions that a false belief would leave the audience in a worse position than ignorance and that the audience's interest is what it appears to be and that is being informed, the best outcome for the audience is to trust when the speaker is trustworthy and not to trust when the speaker is not trustworthy. Call the first outcome the *cooperative outcome*. Whilst, in any particular case, it may lie in the speaker's interest to tell the truth, what explains this lie of interest will be the more basic interest of influencing the audience. A conversation whose ostensible purpose is the giving and receiving of information will be of interest to a speaker because it is a way of getting an audience to believe something and is thereby a way of exerting an influence on the audience. And we have a basic interest in exerting an influence on others. Telling the truth could further this interest but only telling the truth expediently, or telling the truth conditional on one's other interests. Being bound to give an audience the information sought—being trustworthy—would not be in a speaker's interest. So whilst the cooperative outcome would be the best-case outcome for an audience, the commitment of trustworthiness would not be best for a speaker. The best outcome for a speaker would be to receive an audience's trust yet have the liberty to tell the truth or not given the shape of interest in the particular case. However, telling the truth merely when it suits one is a way of being untrustworthy. So the best outcome for a speaker would be the trust and untrustworthiness combination. Since this is the worst case scenario for an audience, the testimonial situation has a structure of pay-offs that resembles the prisoner's dilemma.<sup>4</sup> As such the testimonial situation could be said to present a *problem of trust*: the rational thing for an audience to do seems to be to not trust another for information in the first place.<sup>5</sup> Since the testimonial situation is thereby not one 'in which coincidence of interest predominates', it is not a coordination problem in Lewis's sense. So it seems that Lewis's explanation of social norms cannot be given for the norms of trust and trustworthiness.

5.

Nevertheless, Lewis's explanatory strategy can still be pursued. This strategy is to explain the cooperative outcome in terms of the interests of the parties

135

producing it. Ultimately, what explains a convention is the fact that it is a proper coordination equilibrium. The problem is that in prisoner's dilemma type cases non-cooperation is the dominant option, in that it is the rational thing to do whatever the other party does. However, this all changes given a couple of assumptions about the recurrent situation that presents the dilemma. The first assumption is that individuals have a sufficiently large chance of meeting again so that they care about future interactions, and the number of these future interactions is indefinite. I might depend on you for information on this occasion but you will depend on me on some future occasion and there is the mutual expectation that we will find ourselves in an indefinite number of these testimonial situations. As such there is no final interaction where the trusted party might hope to secure extra benefit by behaving untrustworthily. The second assumption is that each party can make it clear to the other that they will behave in a tit-for-tat manner, which is cooperating at the start but responding to any non-cooperation with retaliation. If the hunter concealed the fact that he made a kill on this occasion and was found out, he would lose the benefit of sharing on another occasion. So whilst non-cooperation could potentially pay dividends on any given occasion, this would be offset by future losses. Given these two assumptions, the cooperative outcome is a coordination equilibrium: no one benefits from unilateral non-cooperation. And it is a proper coordination equilibrium: there is no better strategy. That the cooperative outcome then comes to describe a convention and so be prescribed as a social norm follows on Lewis's account once it is added that everyone expects and prefers conformity. This expectation is delivered by the second assumption: it is clear to all that each is behaving in a tit-for-tat manner and so will conform by default. And the preference for conformity follows given that another's non-cooperation results in loss of the benefits of the cooperative outcome for two rounds, this and the next when retaliation is due. So it is possible to give Lewis's explanation of how the cooperative outcome becomes a matter of convention, where this explanation appeals to little more than rational self-interest.<sup>6</sup>

Philip Pettit (1990: 735) develops a particular implementation of this strategy for explaining norms of cooperation. He starts with three criticisms. First, what this strategy explains is not a norm prescribing a certain behaviour but one prescribing that one behave in this way in a tit-for-tat manner. So it does not offer an explanation of the norm of trustworthiness, or telling the truth reliably, so much as one of the norm of truth-telling in a tit-for-tat way. However, Pettit does not regard this as a significant matter because these norms should be extensionally equivalent in that general tit-for-tat truth-telling would result in the same behaviour as reliable truth-telling. Second, this explanation only works for certain norms. It only works for what Pettit calls 'type B prisoner's dilemmas':

In a type B dilemma, defection by even a single individual plunges at least one cooperator, and perhaps many more, below the baseline of universal defection. In a type A dilemma

136

this is not so and, at limit, the lone defector may have only an imperceptible negative effect on cooperators. (Pettit 1990: 737)

The problem for this explanatory strategy is then that lots of many-party prisoner's dilemmas are type A. However, this is not a problem in this case: the problem of trust is type B because defection by a speaker in the testimonial situation will plunge an audience below the baseline; that is to say, it is worse for the audience to have a false belief than remain ignorant. The real problem, third, is that behaving in a tit-for-tat way requires that people be willing to retaliate. And if the norm being explained is a norm of trustworthiness as opposed to tit-for-tat trustworthiness, then being trustworthy in a tit-for-tat way requires 'that people will break norms punitively in order to punish those who break them for convenience. . . . And this disposition is not genuinely manifested among those who honor norms' (Pettit 1990: 737). Pettit's key observation then follows this third criticism and pursues Lewis's thought that sanction can come by way of others' disapproval. It is that whilst people might not be given to retaliation—such behaviour is costly and risks an escalation of retaliation rather than a return to cooperation—people do retaliate after a fashion: 'a violater can be punished—or of course a conformer rewarded—by the attitudes of others' (Pettit 1990: 739). And 'once these approbative costs and benefits are put into the equation', Pettit hypothesizes, 'we can see our way to explaining why the emergence and persistence of otherwise puzzling norms may be unsurprising' (1990: 742). That is, once these costs and benefits are put into the equation we can make good the tit-for-tat implementation of Lewis, (except we now understand 'tating' as holding a disapproving attitude). What is assumed here is that people care about others' approval and disapproval. If it is also assumed that these feelings are automatic and costless, then one gets sanctions for free and one has an explanation of compliance with the norm: the sanction which ensures compliance is not the threat of retaliation but the fear of disapproval and desire for approval. So we can regard trustworthiness as the best preference given trust, and the cooperative outcome as a coordination equilibrium given the approval it elicits. Thus starting with Lewis's suggestion that conventions are norms because they specify regularities which everyone recognizes that one ought to conform to. And making it explicit that the feeling *that one ought to conform* underwrites approval and disapproval and plays a causal sustaining role gives Pettit's Lewisian account of social norms:

A regularity, R, in the behaviour of members of a population, P, when they are agents in a recurrent situation, S, is a *norm* if and only if, in any instance S among members of P,

- (1) nearly everyone conforms to R;
- (2) nearly everyone approves of nearly anyone else's conforming and disapproves of

- (1) nearly everyone conforms to R;  
nearly anyone else's deviating; and
- (3) the fact that nearly everyone approves and disapproves on this pattern helps to ensure that nearly everyone conforms. (Pettit 1990: 731)

137

In short, norms are regularities in behaviour that we approve of where this approval and corresponding disapproval enforce the regularity. The approval and disapproval are rationally intelligible because the norm specifies the cooperative outcome in some situation of collective action. And this outcome can be seen to be a proper coordination equilibrium once our concern for approval is factored into the equation.

6.

Pettit distinguishes two different ways of explaining norms of cooperation. 'Behavioural explanations' show how certain patterns of behaviour are a matter of self-interest. The tit-for-tat explanation of cooperation is behavioural; it shows how cooperation can 'emerge in a world of egoists without central authority' (Axelrod 1984: 3). Pettit contrasts his 'attitude-based explanation' which shows 'first why certain attitudes of approval are intelligible . . . [before] then showing how they might generate the patterns of behaviour required for norms' (Pettit 1990: 733). Now, suppose that agents *S* and *A* are in a joint action situation *T* and there is a norm stating that *S* should  $\phi$  in *T*. Pettit's account as to why disapproval of breach of a norm like this is intelligible is that the norm prescribes a cooperative outcome that is a proper coordination equilibrium. So were *S* not to  $\phi$ , *A* would disapprove of *S*'s failure to  $\phi$  because it is to his, *A*'s, detriment, and *A* thinks that *S* has a reason to  $\phi$ . This renders *A*'s attitude of disapproval intelligible after a fashion: it renders it intelligible from the perspective of what lies in *A*'s rational interest. However, this is not intelligible from *A*'s perspective in that it does not detail *A*'s reason for his attitude of disapproval. The judgement of wrongness made when a norm is breached does not consist of a judgement of consequence: *A*'s disapproval does not stem from the loss that will follow from *S* not  $\phi$ -ing. Rather, *A*'s judgement is no more than that *S* should have  $\phi$ -ed, or should have  $\phi$ -ed given that he was in situation *T*. Consider the question, what exactly do people disapprove of when they disapprove of others breaking a norm? Or, what makes disapproval intelligible to those making a disapproving judgement? These questions will be answered by appeal to the norm: the norm that one should  $\phi$  in *T* describes a standard of conduct that is appealed to when making an approving or disapproving judgement. What this means is that one cannot see the attitudes of approval and disapproval as intelligible independently of acceptance of the norm that articulates these judgements and that is appealed to in order to justify them. Consequently, these judgements do not have the kind of independent intelligibility necessary for Pettit's 'attitude-based' explanation. That is, *unless* intelligibility is really just a matter of rational self-interest. But then patterns of behaviour are ultimately approved of because they are

in everyone's best interest, and the distinction Pettit draws between styles of explanation is not substantive. But this is just to say

138

something that Pettit (1990: 725) otherwise acknowledges, which is that his account is game-theoretical.

On Pettit's definition of social norms, general attitudes of approval and disapproval 'help to ensure' conformity to the regularity in behaviour that is the outward part of the social norm. This is undoubtedly true, but the question is what is the right reading of 'helps to ensure'. The right reading for Pettit's 'attitude-based' explanation is a strong one like 'is the cause of'. The idea is that conformity is explained once people's preferences are suitably adjusted to take into consideration the fact that we care about others' opinions of us. The problem for this explanation starts when we take seriously the question, what exactly do people disapprove of when they disapprove of others' breaking a norm? In terms of the schematic scenario just described, what *A* disapproves of when *A* disapproves of *S* not  $\phi$ -ing in *T* is just the fact that *S* didn't  $\phi$  when *S* should have done. The norm itself articulates *A*'s reason for disapproval. Social norms are thereby held as standards of behaviour that justify both attitudes of approval and disapproval. To use a sociological term of art, people *internalize* social norms, which is to say that social norms describe patterns of behaviour that people are motivated to follow for no other reason than that these patterns of behaviour are valued in themselves or held as ultimate ends. Let me describe this kind of motivation as *intrinsic*.<sup>7</sup> The idea that we internalize social norms is the idea that we are intrinsically motivated to behave in the way the norms prescribe. This aspect of social norms, the fact as Elster says that they 'have a *strong grip on the mind*', is not represented by game theory, which is not concerned with intelligibility from the inside or intelligibility in terms of the norm. So returning to the issue of *A*'s disapproval of *S* not  $\phi$ -ing, this disapproval is the kind that is meant to explain, on Pettit's account, general conformity with the norm to  $\phi$  in action situation *T*. But if *A*'s disapproval stems from the fact that *A* has internalized this norm then it is this fact and not others' disapproval that should explain *A*'s compliance. But what goes for *A* goes for *S*; these are just agent role placeholders. This is to say that whilst others' disapproval can cause conformity, the reason for others' disapproval will also centrally be their reason for conformity. Returning to Pettit's criticisms of the tit-for-tat behavioural explanation of social norms, these criticisms can now be seen to be more substantial than Pettit acknowledged. If peoples' reason for conformity is their having internalized the norm they conform to, then we have an explanation as to why people tend not to break the norm punitively. To break the norm punitively is still to break the norm and so fall short of the standard of behaviour prescribed. And if the norm articulates peoples' reason for conformity, then it matters whether the norm is to  $\phi$  or to  $\phi$  in a tit-for-tat way. Even if these two norms are extensionally equivalent, they give different reasons for action. So the tit-for-tat explanatory strategy is limited from the start.

The idea that social norms articulate reasons for action and figure in justifications of attitudes of approval and disapproval allows for a very simple explanation as to why people conform to social norms. We obey social norms because these norms describe what we believe that we ought to do. We obey social norms because we have internalized them. This is more properly an ‘attitude-based’ explanation of social norms, and can be labelled the ‘social’ account (since talk of internalizing social norms is commonplace in the social sciences). However, whilst it provides a very simple answer to the question of why people conform, the social account raises two further questions. First, we might have a rudimentary explanation of why people comply with social norms, but can this be filled out? Second, game-theoretical or Lewisian accounts of social norms make good sense of why norms exist in that they provide an explanation not merely of compliance but also of the norms themselves. So one might also hope for an answer to this question: why do particular norms exist? In the next section I consider this question, and its particular form: why is it that we have social norms of trust and trustworthiness?

7.

In *Truth and Truthfulness* Bernard Williams offers an imaginary genealogical account of what he calls the *virtues of truth: Accuracy and Sincerity*. These are the dispositions to care about the truth of one's beliefs, and to come out with what one believes. Williams's genealogy offers an explanation of *our valuing* these dispositions or virtues of truth. It is because we value these dispositions that we try to get things right in belief and utterance. However, our valuing Sincerity is not a matter of the crude prescription: ‘always tell the truth’. The mail opener, referred to above, does not manifest the disposition of Sincerity in Williams's sense, even though she is sincere in believing what she says. When the conversation is one of giving and receiving information, Sincerity is a matter of being appropriately informative: it is a matter following the Cooperative Principle and its maxims, which define the social norm of trustworthiness. So Williams's genealogical justification of Sincerity offers a way of explaining why it is that we have this social norm, and the paired social norm of trust. Moreover, Williams's genealogical justification focuses on the joint action situation that is the testimonial situation which presents the problem of trust.<sup>8</sup>

Williams's genealogy starts by imagining a State of Nature consisting of a primitive social group with limited technology and no writing. Although primitive, this social group is imagined to be a real society whose members have projects and interests, and are related to one another in various ways and via various roles. As with any society, the society imagined in the State of Nature will

involve cooperative engagements which demand information be communicated between individuals. Given that an individual can only be at one place at one time, individuals will often gain what Williams (2002: 42) calls *purely positional advantage*; that is, by virtue

of their location at a time, one individual can come to possess information that another individual needs. It follows that even in the State of Nature, thus minimally characterized, Accuracy and Sincerity are desirable from the social point of view; they will be socially valued because pooled information is a social good and necessary for many cooperative endeavours. However, possessing the disposition of Sincerity need not always be in an individual's best interest. Williams gives the example of the hunter who has found prey that he would rather keep for himself and his family. This raises the problem that:

The value that attaches to any given person's having this disposition [Sincerity] seems, so far as we have gone, largely a value for other people. It may obviously be useful for an individual to have the benefits of other people's correct information, and not useful to him that they should benefit of his. So this is a classic example of the 'free-rider' situation.

(Williams 2002: 58)

The problem is that the collective valuing of Sincerity does not itself give an individual a reason to value Sincerity, or be sincere. Whilst it is always in an audience's interest to be informed, Sincerity needn't best serve a speaker's interest and as audiences we know that this is the case. The problem that Williams thus identifies is the problem of trust; and since the State of Nature represents a basic society, the possibility that a conversation as to the facts could be stymied by this problem shows 'that no society can get by . . . with a purely instrumental conception of the values of truth' (Williams 2002: 59).

What any society requires is that individuals have internalized Sincerity as a disposition, where this is to say that individuals are motivated to act in a sincere way simply by the description of this way of acting as sincere. Where this is true, Sincerity will have *intrinsic value* (or intrinsic value in the society). Something's having intrinsic value, Williams then goes on to argue, can be understood in terms of the satisfaction of two conditions. For something, X say, to have intrinsic value in a society: first X must be 'necessary, or nearly necessary for basic human purposes and needs'; and second X must 'make sense to them [the society members] from the inside, so to speak' (Williams 2002: 92). The first of these desiderata is established by the imagined genealogy. If Sincerity is not given intrinsic value, then any conversation that purports to be one of giving and receiving information will generate the problem of trust. This threatens to stymie both the conversation and any further cooperation. However, we do cooperate in conversations as to the facts. We tell one another what we know and we have a way of life wherein testimony is a source of knowledge. So whatever needs to be in place to avoid the problem of trust must be in place and that is that we intrinsically value Sincerity. We must be motivated to be sincere as an end in

itself. The second desiderata is then giving an account of how this motivation is made sense of.

Since any value is made sense of through its connection to further values, how a society gives intrinsic value to Sincerity can be philosophically unearthed through conceptual

analysis. What conceptual analysis shows about *our* social history is that *we* understand Sincerity through its relation to trust and our valuing trustworthy behaviour. *Sincerity is trustworthiness in speech*. This is more than the avoidance of lying; our being trustworthy can require our lying. Equally, it is not simply the disposition to say what one believes; one can implicate falsehoods by saying what one believes and so be untrustworthy by doing this. The mail opener implicates that someone else has been opening your mail. What Grice's discussion of implicature then shows is that:

Implicature do not presuppose language as simply a practice involving semantic and syntactic rules, together with the norm that certain kinds of utterances are taken to be true; they look to the use of language under favourable social conditions which enable it to be indeed co-operative. They are *conversational* implicatures, but not everyone who is talking with someone else is engaged, in the required sense, in a conversation. What is required for that to be so are certain understood levels of trust. (Williams 2002: 100)

We have achieved these levels of trust because we intrinsically value trustworthiness in speech, which is the disposition of Sincerity. And this is just to say, I suggest, that the norm of trustworthiness, which is the prescription that speakers follow the Cooperative Principle and its maxims, is internalized as a social norm. In learning to have conversations one learns this norm, and the presumption that things are as the norm prescribes then allows us to uncover implicatures or what people mean by what they say. Since we can tell people what we know by implication as much as by bald statement, our norm of trustworthiness is then necessary for testimony being the source of knowledge that is. The genealogical justification that Williams offers in *Truth and Truthfulness* for our having the disposition of Sincerity can then be presented as a genealogical justification of the norm of trustworthiness, and with it the paired norm of trust. This addresses the challenge of explaining why we have these particular norms.

8.

The problem, as Williams is aware, is that the claim that X has intrinsic value faces a dilemma. Left like this it is mysterious as a claim. Why should the description of an act as X be a motivation to act this way? But if the mystery is explicated, then the account of the value threatens to become reductive with X being merely instrumentally valuable. Williams's two-condition account is meant to address this dilemma; genealogy is meant to achieve 'explanation without reduction' (2002: 90). A similar problem faces social accounts of social norms. The explanation of

behaviour that states that people act a certain way because they have internalized a certain norm is not fully satisfactory issuing in the question: but why should they be motivated to act in this way? This is the first challenge noted in section six, and here game-theoretical, or Lewisian, accounts of social norms seem to be genuinely explanatory since they have self-interest as the final appeal, which does not seem to just raise further questions. The problem with this, I argued, is that if social norms are



rationalized in terms of rational self-interest, then what is left out is the sense that we find in acting in the way that the norm prescribes. What is left out is our understanding of the value that motivates our action. However, to make good the claim that this genuinely impoverishes these explanations, what is needed is a fuller account of *how* our understanding of value motivates action. Moreover, the need to meet this challenge is sharp in this particular case because, on the face of it at least, we do find trust problematic. That this is so can be easily reinforced by considering a testimonial situation but stripping away all the factors that could be used in a game-theoretical solution to the problem of trust. One must imagine that there is no sanction on untrustworthy behaviour and that the speaker does not care for the good opinion of others, or at least that the audience cannot have the assurance of believing either of these things. That is, one must imagine a case where an audience is engaged in a conversation as to the facts with a speaker whose particular motivations and preferences the audience is ignorant of. In this case, the worry that the speaker will not tell the truth, or will only do so if it suits them is a natural worry. Thus the idea that it is reasonable to trust ‘if we know absolutely nothing about someone’, Williams describes as simply ‘a bad piece of advice’ (2002: 111). He then adds that

[i]t may be said that a hearer never has a reason for believing that P which lies just in the fact that a given speaker has told him that P. He has to believe also that the speaker (on such matters, and so on) is a reliable informant. (Williams 2002: 77–8)

So our intrinsically valuing Sincerity is *not sufficient*, according to Williams, to ground reasonable acceptance testimony as to the facts. What is also needed is the belief that a bit of testimony is reliable or that a speaker is manifesting the disposition of Accuracy. Now, I think that this is the wrong way to go, and the wrong way for Williams to go: the attitude of trust can suffice for reasonable acceptance of testimony. And I think that this is what is delivered by considering how our valuation of Sincerity ‘makes sense to us from the inside’. However, leaving this argument until the next section, the point to be made here is that the temptation to require more than trust for reasonable belief is an expression of finding the problem of trust genuinely problematic. So if this problem of trust is meant to be solved by our giving intrinsic value to Sincerity, then a fuller statement is needed as to how this locus of value motivates our acting in certain ways. This is the first challenge facing social accounts: filling out how the internalization of social norms explains conformity with them.

What is needed, in the case of the norms of trust, is recognition of how trust figures in our explanations and justification of action. These explanations and justifications are rather straightforward. Asked why we took the risk of depending on someone we often answer simply that we trusted them, and asked why we put ourselves out to do something we answer that someone trusted us to do this thing. Suppose that one person A trusted another S to do something, to  $\phi$ . And suppose that A trusted S to  $\phi$  in the thick sense that A depended on S  $\phi$ -ing and expected this to be at least part of S's reason for  $\phi$ -ing.<sup>9</sup> Such an

attitude of trust is quite common and might be found in the coordination problem described in section three. In this case, *A* trusts *S* to turn up at the The Lamb at eight and *A* trusts *S* to do this in the sense that *A* thinks that at least part of *S*'s reason for turning up at this pub at this time is the fact that *A* depends on *S*'s doing so. In this case if *S* were asked why he was going to this pub at this hour he might reply 'to meet *A*' or '*A*'s waiting for me' and if pushed to take another course of action, *S* could emphasize the reason this gives by making it explicit: '*A* trusts me to turn up'.<sup>10</sup> So we can use the fact that another has trusted us to do something to explain why we did this thing. And we can use the fact that we trust someone to explain why we showed a willingness to depend on them in certain ways: asked why he was bothering to get to the pub on time, *A* might reply that he trusts *S* to show up at this time. So the attitude of trust, in this thick sense, figures in justificatory explanations of action. This, I suggest, is what the claim about intrinsic value amounts to: we credit these kinds of ways of making sense of things. Since we use the attitude of trust to explain and justify acts of trust and trustworthiness, our having the dispositions that follow from internalizing norms of trust and trustworthiness 'makes sense to us from the inside'. However, to say that we use the terms of these norms to explain and justify, is to say that we are motivated to act in the same terms. This is what the idea of internalization delivers: the prescription of the norm captures the way the subject thinks about the action prescribed. The idea that we have internalized the norms of trust and trustworthiness then offers a genuinely explanatory account of action because if it is true that we have these dispositions to trust and be trustworthy, then the prescriptions these norms make will outline good descriptions of our reasons for acting.

9.

Williams's official solution to the problem of trust found in the State of Nature is that this problem is resolved by finding some way to give Sincerity intrinsic value.

144

We then give Sincerity intrinsic value by taking it to be a form of trustworthiness and valuing trust and relations structured by trust. If this claim about how we give Sincerity intrinsic value is understood as a claim about our finding certain descriptions and justifications of action persuasive, then the shape of the solution to the problem of trust Williams's genealogy offers is revealed.

The problem of trust presupposes an account of the kinds of reasons people have for acting. On this account, action is explained in terms of the agents' beliefs and desires. Epistemic rationality demands that an audience desires to believe the truth and avoid falsehood and so has a preference for ignorance over error. Then absent any belief about a speaker's motivations, or any grounds for predicting the probable truth of utterance, and the result can only be that it is not reasonable for the audience to accept what he is told. Game-theoretical solutions to this problem add grounds for belief. A concern for others' good opinion in Pettit's case. With the problem situation then reconfigured, a coordination equilibrium is found and the norms of trust and trustworthiness emerge. Williams recognizes the extent of the problem, offers the basis of a solution in terms of

giving intrinsic value to Sincerity, but then feels compelled to add that some belief about the truth of utterance is still necessary. However, his idea that Sincerity be intrinsically valued suggests an alternative solution. On this solution, what goes wrong with the problem of trust, why trust is seen as problematic, is that the only explanation of acceptance allowed is one that proceeds in terms of an audience's belief and desires. However, we can act out of trust: our trusting a speaker for the truth can give us sufficient reason to accept what the speaker tells us. An audience can explain why he accepted what a speaker told him, and so detail his motivations and justify his acceptance, by saying that he trusted the speaker for the truth. However, an explanation of action that is couched in terms of trust cannot be translated into one couched in terms of belief and desire. This is because the attitude of trusting someone to do something involves placing an expectation on that person: that they will act in a certain way and for a certain reason. And this *expectation of* them is not the *expectation that* something will happen. The difference between these kinds of expectation is marked by the fact that when we expect things of people we are susceptible to various reactive attitudes if they do not act as we expect. So given a conversation as to the facts and an audience who trusted the speaker for the truth and was misled, this audience will be liable to resent the speaker's actions. The susceptibility to such a feeling of resentment defines the expectation as one that is placed on a person or held of them and distinguishes it from the expectation or belief that something will happen. And this feeling of resentment involves commitment to the norm of trustworthiness as an objective standard: any resentment felt will not be mollified by the knowledge that the trusted individual had no inclination to tell the truth because what is felt is that the trusted individual *did have* such a reason and *should have* acted on this reason. This is the reason described by the norm of trustworthiness, which is meant to prescribe behaviour irrespective of subjective

145

motivation or personal interest. So the norms of trust and trustworthiness define standards of behaviour within which explanations in terms of trust make sense and are accordingly good explanations. This is what is missed when trust is seen as problematic.

Where norms of trust and trustworthiness are internalized, the social background will be one of 'certain understood levels of trust'. It will be such that if an audience *A* trusts a speaker *S* for information, this will give *S* a reason to tell *A* what he needs to know; and if *S* tells *A* something, this gives *A* reason to accept what *S* tells him. The idea that conversation can be structured by presumptions of trust then allows for the following explanation of what goes on, or should go on, in a conversation as to the facts. A speaker's reason for telling an audience what he does—the explanation of the speaker's testimony—will be the speaker's perception that the audience depends on him for this information. In this case, if the speaker *S* tells the audience *A* that *p*, it will be because *S* believes himself to know that *p*, and assumes responsibility for letting *A* know that *p* in the following sense: *S* takes it on himself to tell *A* that *p* if and only if he, *S*, knows that *p*. In this way *S* is trustworthy. And in recognizing that the speaker *S* intends that *A* come to believe that *p* and trusting *S*, *A* will then take *S*'s telling him that *p* as something like a

promise that  $p$  is true. Testimony then functions to transmit knowledge from speaker  $S$  to audience  $A$  because it transmits the responsibility for justifying belief from audience  $A$  to speaker  $S$ . This explanation is offered by Richard Moran (2006) and following his lead call it the *assurance* explanation. On this explanation  $A$  doesn't need the belief that  $S$  is reliable to have a reason to accept what  $S$  tells him. Rather,  $A$ 's trusting  $S$  gives him a reason through delivering the presumption that  $S$  is trustworthy. Moreover, for  $A$  to seek further reason to believe that  $p$  would be to reject  $S$ 's assurance that  $p$  is true. This would be as likely to provoke  $S$ 's resentment as straight disbelief since, Moran argues, it amounts to a refusal to accept  $S$ 's assumption of responsibility in telling  $A$  that  $p$  (see Moran 2006: 301). And this feeling of resentment, I suggest, is parallel to that a misled audience would experience in that it equally involves commitment to a social norm as an objective standard, in this case the norm of trust. Supposing  $A$  did seek the support that  $S$  was reliable prior to belief,  $S$ 's resentment at not being believed would not be appeased by the knowledge that  $A$  had no inclination to belief because what is felt is that  $A$  did have a reason for belief, given by  $S$ 's telling, and should have believed for this reason. So violations of either norm of trust—disbelieving a speaker or misleading an audience—will engender resentment and other punitive attitudes. So let me drop the demand for a belief about Accuracy or reliability from Williams's genealogy: it is better to see the problem of trust confronted in the State of Nature as resolved by the establishment of levels of trust that allow the giving and receipt of testimony as assurance.

How does the existence of social norms of trust and trustworthiness bear on epistemological theories of testimony? It is now possible to give a brief

146

answer to this question. One implication is that non-reductive theories are correct to describe our attitude towards what others tell us as trusting. However, non-reductive theories, I suggest, are wrong then to hypothesize that we are default justified in trust; the norm of trust is a social norm and not a general or universal epistemological principle. Whilst every society confronts the problem of trust, since it is confronted in the State of Nature, securing the necessary motivations through a valuation of trust is but one solution to this problem. Other norms are possible. Another implication is that insofar as these social norms do operate in our society, there should be plenty of scope to give a reductive theory of testimony: we should have good evidence that tellings will prove generally reliable even if evidence of particular reliability is hard to come by.<sup>11</sup> However, the possibility of this defence of reductive theory should not constitute a justification of this position because the norm of trustworthiness is associated firstly with testimony being a source of knowledge that can be explained in assurance terms. And this, I think, is the central implication. Our having the social norms of trust and trustworthiness is a function of our having a way of life wherein we have conversations as to the facts and tell one another what we know. The epistemology of testimony cannot be synonymous with the epistemology of tellings: there are too many messy and varied cases, which determine that the assurance can only be part of the story. But if we have these norms of trust and

trustworthiness, then the assurance explanation of how another's telling can put us in a position to know something must be an essential part of our epistemological story.<sup>12</sup>

## 8 Testimonial Entitlement and the Function of Comprehension

Peter J. Graham

Why do you think it will rain tomorrow? Because, you say, it always rains this time of year; you can see dark clouds for miles; and besides, you read the weather report. The reasons you offer *justify* your belief; they comprise a *justification* for your prediction about the weather. Justifications are commonplace. Justifications involve reasons or evidence you often cite when asked how you know. Justifications support your ability to *justify*—to argue in favour of—your belief, to show that your belief is (likely to be) true (Unger 1968; Audi 1988; Kornblith 2008).

Justification is one kind of positive epistemic status or standing. Positive epistemic standings are goods, successes, fulfilments, or achievements understood in terms of promoting true belief and avoiding error (Alston 1985, 2005).

Epistemic justification is not the only kind of positive epistemic standing. Perceptual beliefs paradigmatically enjoy positive epistemic status, but higher non-human animals, human infants, children, and ordinary adults often lack evidence or reasons in support of their perceptual beliefs, and even more often, if not always, lack the ability to cite evidence or reasons in support of their perceptual beliefs. Ask a small child why he believes the ball continues to exist as it passes behind the screen, and at best you'll get a gurgle or a burp. Ask a chimpanzee trying to open a crate why she believes there are bananas inside, and you won't get any answer at all. Well-formed, garden-variety perceptual beliefs enjoy positive epistemic status, but that status does not depend on reasons or evidence, and it especially does not depend on the subject's ability to offer reasons or evidence (Goldman 1979).

I use the term 'warrant' for positive epistemic statuses that provide epistemic support or grounding in favour of holding the belief. Justifications are a familiar kind. Evidentialism arguably captures the contours of this kind of warrant. But justification is clearly not the only kind of warrant. I use 'entitlement' for positive epistemic support or grounds in favour of belief when the support or grounding

148

does not involve or depend upon evidence or reasons, especially evidence or reasons the subject might cite in favour of his or her belief (cf. Burge 1993, 1996). Entitlements are goods, successes, or fulfilments in the way of promoting true belief and avoiding error that do not depend on evidence or reasons available to the subject. Though many perceptual beliefs enjoy justifications, paradigmatically they enjoy entitlements.

What about so-called 'testimony-based beliefs?' These are beliefs based on our capacity to comprehend assertive speech acts. What kind of warrant do they enjoy? When you tell

me the rain won't let up and I comprehend and accept your assertion, what kind of warrant does my comprehension-based belief enjoy, justification or entitlement?<sup>1</sup>

On the traditional view, warrants for comprehension-based beliefs rely exclusively on entirely first-hand justifications. This tradition has fallen on hard times. For though comprehension-based beliefs are by and large warranted, it is now widely thought that fully adequate, entirely first-hand justifications are in short supply (Coady 1992; Fricker 1995).

Though comprehension-based beliefs sometimes enjoy fully adequate, entirely first-hand justifications, and very often enjoy partially adequate, partially first-hand justifications (cf. Faulkner 2002), I shall argue that comprehension-based beliefs formed though a process of filtering generally enjoy *prima facie pro tanto* entitlement.<sup>2</sup> We thus need not worry over the paucity of fully adequate, entirely first-hand justifications. Though justifications clearly play an important role in a complete account of the epistemology of comprehension-based beliefs, they do not play the *only* role.

Why should beliefs based on comprehension-via-filtering enjoy entitlement? I shall first argue for an account of (a kind of) entitlement and then apply it to comprehension-based beliefs. A belief enjoys (a kind of) *prima facie pro tanto* entitlement when based on a normally functioning belief-forming process that has

149

forming and sustaining true beliefs reliably as a function.<sup>3</sup> And comprehension-with-filtering has that function. Beliefs based on comprehension-with-filtering thus enjoy (a kind of) entitlement when our capacity to comprehend and filter functions normally, for our capacity has forming and sustaining true beliefs reliably as a function. When functioning normally, comprehension-with-filtering confers *prima facie pro tanto* entitlement on the beliefs it causes and sustains.

In the first section I say a few words about our psychological capacity to comprehend and filter. In the second I explicate function, normal conditions, and normal functioning. In the third I put these notions to use to argue for my first main thesis, that (a kind of) epistemic entitlement derives from the normal functioning of the belief-forming process when the process has forming true beliefs reliably as a function. In the fourth I argue indirectly for my second main thesis by first arguing that *assertion* has causing and sustaining *true* hearer beliefs as a function. In the fifth I defend this claim against four counter-arguments. I then argue directly in the sixth for my second thesis that *comprehension-with-filtering* has causing and sustaining true beliefs *reliably* as a function. Though I have tried to be concise, the paper is not short enough to be easily read without taking a break. A natural place to pause would be at the end of section three. Taking a moment to digest the first main thesis may help prepare the palette for the second. Taking a break after the fifth section might not be a bad idea either.

## 1. COMPREHENSION-WITH-FILTERING

For the most part, we automatically, quickly, and reliably comprehend speech. Comprehension involves a number of factors: standing abilities to comprehend shared meanings of words, shared awareness of relevant features of the context of utterance, and standing expectations of conformity to cooperative principles or maxims governing speech behaviour, not to mention the rest of the language suite: syntax, morphology, phonology, and speech perception.

150

Usually we correctly comprehend what people say. But sometimes we make mistakes: we sometimes miscomprehend and thereby misrepresent the content or force of a speech act, or both.

Comprehension states are exercises of our capacity to comprehend. Since my focus is assertive communication, I shall use ‘comprehension state’ throughout for comprehension states that purport to represent assertions.

I use ‘assertion’ broadly to cover the whole range of constative speech acts: affirm, allege, avow, certify, concur, declare, describe, disclose, inform, predict, recount, report, say, state, submit, tell, and testify, among others.<sup>4</sup>

Comprehension-based beliefs are beliefs caused or sustained by taking comprehension states at face value. In the ordinary case, a speaker (sender) asserts that *P*. The hearer (receiver) represents *P* as asserted.<sup>5</sup> The representation of the speech act—the hearer’s comprehension state—disposes the hearer to believe that *P*. In the default case, comprehension prompts belief in the propositional content of the assertion as represented. Being so moved constitutes taking comprehension at face value. My topic is the epistemology of beliefs so formed. I thus use ‘testimonial entitlement’ to refer to the entitlement that beliefs based on comprehension enjoy.

151

Our developmental starting point is not a broadly Humean, sceptical stance, where we withhold belief until sufficiently strong supporting reasons come along. Rather our developmental starting point is a broadly Reidian, credulous stance (Harris 2002; Koenig and Harris 2005). Very young children nearly automatically accept whatever they comprehend. We are built to take comprehension states at face value; credulity is the default.<sup>6</sup>

That said, we obviously do not believe everything we hear or read. Filtering soon develops. If you label something right before a child’s eyes when she already confidently believes the opposite, she is less apt to believe you. If you call a cup a ‘dog’ right before a 16-month-old, the child will direct more attention to you than she will to speakers who correctly label the object. If you do the same thing before an 18-month-old, she may even correct you (Koenig, Clement, and Harris 2004).



Another filter involves manner of presentation. Three- and four-year-olds are sensitive to how confident an assertion is. An unsure, less committal tone of voice sometimes makes acceptance less likely. Conversely, more confident assertions make acceptance more likely (Moore, Bryant, and Furrow 1989; Jaswal 2004).

Young children acquire track record information, either first hand or from others. On the basis of this information, they tend to selectively accept assertions. When children have information that a speaker is reliable, that speaker tends to prompt belief more often than the speakers which children 'know' are less reliable. Again, it is three- and four-year-olds that display this ability, with four-year-olds showing more sensitivity to such information (Koenig, Clement, and Harris 2004). This stands to reason, for only after children acquire the concepts of belief and truth are they in a position to form beliefs about which of their interlocutors are reliable, beliefs they can then rely upon in blocking or supporting the transition from comprehension to belief. (I have in mind the false belief task.)<sup>7</sup> Coherence-checking, a fourth form of filtering, emerges later. I will have more to say about coherence checking further along. As we mature, we rely on these filters and more.

152

Filtering need not involve explicit awareness—belief or judgement—that counter-considerations are absent. Filtering need not involve reasons and reasoning. The subject need only respond by dampening, if not entirely suspending, the normal force of comprehension (Goldberg and Henderson 2006).

In sum, filtering involves sensitivity to counter-considerations: were there are counter-considerations of a certain sort, acceptance would be less likely. Different filters involve different sensitivities. We obviously do not believe whatever we are told; credulity as default does not imply unalterable gullibility. We're no fools.

Sometimes filtering blocks the normal force of comprehension, but we still come to believe what we were told. Consider Nick and Oscar. Nick correctly comprehends Oscar's assertion. But Nick is suspicious of Oscar's credibility. Perhaps Nick antecedently doubts that Oscar will tell the truth. Perhaps Oscar's story does not fit with other things Nick believes. Perhaps Oscar's report contains internal contradictions. Or perhaps Oscar is alarmingly fidgety. If this were the end of the matter, Nick would not believe what Oscar just said. But on reflection Nick thinks Oscar *can* be relied upon. Taking as a premise that Oscar asserted that *P*, Nick comes to believe that *P*. Nick's warrant is a justification, an argument from Oscar's assertion that *P* as premise (and other believed propositions) to the conclusion that *P*. Nick's belief is, as it were, *from* comprehension without being *through* comprehension (cf. Audi 1997).

Once we mature, comprehension involves a 'default-trigger-and-evaluate' structure. Peter Lipton explains:

[I]n most contexts the hearer simply accepts what is [taken as] said, without engaging in any conscious evaluation or inference . . . [But] there are diverse conditions that may trigger the hearer to switch into evaluative mode, where he pauses to consider whether he

should believe what he has been told. . . . The obvious triggers include cases where what is claimed obviously contradicts firmly held beliefs, cases of contradictory testimony, and cases where there is reason to worry about incompetence or insincerity. (2007: 241)

Comprehension states dispose belief, but not willy-nilly. Though credulous, we do not believe everything we are told. We filter and, on occasion, evaluate.

## 2. ETIOLOGICAL FUNCTIONS

I now turn from psychology to teleology. In this section I explicate the notions of function, normal functioning, and normal conditions. There are at least three kinds of functions. My account of entitlement employs the third.

First, there is the familiar category of intended (consciously assigned) functions for items we design to produce certain effects. Artefacts (tools) have many of their functions because of what we have in mind when making them. Usually an artefact's function is what it was consciously and intentionally designed to do. Belts were designed to hold up trousers. Watches were designed to tell the

153

time. Cars were built to take us from point A to point B. The artefact was made, shaped, arranged, composed, manufactured, in order to produce a certain effect. The effect is what the designer had in mind that explains why she made, shaped, arranged, composed, or manufactured the artefact. Call these d-functions.

Second, there are functions fixed by one-off, occasion specific intentions. Intentions are goals, aims, or purposes. Like function, intention is a teleological notion. An intention on an occasion explains why someone does something, or what he or she is using an item for. I went out to get a breath of fresh air. I put the stapler in the doorway to keep the door ajar. Call these functions i-functions.

A third kind of function applies paradigmatically to biological organs and systems, though it has a much wider application. It even applies to some functions of artefacts; not all functions of artefacts are d-functions. I have in mind etiological functions, where the function of an item essentially depends on certain explanatory features of its history. Call etiological functions e-functions.<sup>8</sup>

An e-function of an item is the effect of ancestors of the item that explains why the item was replicated, and so why the item exists now, where the item (or the system of which it is a part [Buller 1998, 2002]), underwent some form of selection (natural, artificial, cultural, or some other form). Take the heart. Its function is to pump blood. Hearts now exist because ancestors of hearts pumped blood, and in so doing contributed to survival long enough for new hearts. The eye is for seeing. Eyes now exist because ancestors provided information about the spatial array of objects in the environment. For most animals, sexual reproduction is the mechanism of copying where later hearts are produced on the model of past hearts. But other methods of copying will do. E-functions arise when certain effects of ancestors of an item explain why the item persists.<sup>9</sup>

E-functions are real features of biological kinds. When evolutionary biologists taxonomize according to function, they look for effects that played an important role in explaining why the organs, systems, or activities of animals persist.

But as I have already said, e-functions are not simply biological functions. E-functions are effects of ancestors of any item that played a role in explaining why reproductions of that item exist, no matter what the item is, or

154

how it was reproduced. Learned behaviours, linguistic devices, conventions, cultural practices, games, and so on, may all have etiological functions (Millikan 1984).

The logic of e-functions comes with an informative account of what counts as normal conditions and normal functioning (Millikan 1984: 33–4). When an item has an etiological function, there is an explanation in terms of that function for why the item exists. It exists now, in part, because its ancestors, when working or operating a certain way in certain circumstances, produced the effect that contributed to the item being replicated. That effect becomes the item's etiological function. The way the item operated according to that explanation counts as *normal functioning*. If the current item works or operates that way (or in relevantly similar ways; there are borderline cases), then it functions normally; it operates the way it is supposed to operate. And the conditions according to that explanation (or conditions of similar type; there are borderline cases), fix *normal conditions*. Take the heart. It beat (worked, functioned) a particular way in an organism's body in a particular environment. By beating that way, it pumped blood and thereby contributed to the survival and reproduction of the organism. Pumping blood thereby became the function of the heart. Normal functioning is then just beating or working the way it did that contributes to the explanation of the heart's acquiring the function of pumping blood.<sup>10</sup>

This explanatory, individuating interconnection (between function ascription, function fulfilment, normal conditions, and normal functioning) means that for any item with an e-function, *ceteris paribus* it will fulfil its function, often enough, in normal conditions when functioning (operating, working) normally. If it did not, then the item would not have been replicated because of that effect; it would not have that e-function. This falls out of the logic of e-functions.

155

This fact marks an important difference between e-functions, on the one hand, and d-functions and i-functions, on the other. You can design something to produce a certain effect, but it may fail to ever do it. Think of the dustbin of failed inventions. You can also intend to do something but not succeed. If only life were so easy. Not so for etiological functions. No type of item with an etiological function can have tokens that fail across the board.

We now have an important explanatory fact to rely upon for items with etiological functions: *ceteris paribus* normal functioning in normal conditions leads to function

fulfilment. We can then partially *individuate* and *explain* normal functioning in terms of normal conditions and function fulfilment. If you know the etiological function of an item, functioning normally *just is* working or operating in the way that, according to the historical, normal explanation for the item, fulfils that function often enough *ceteris paribus* in normal conditions. Normal functioning is thus individuated and understood in terms of function fulfilment.

You can fulfil your function without functioning normally, and function normally without fulfilling your function. Stimulate a diseased heart artificially and it may pump blood, though not because it is working the way it is supposed to. You can also take a heart out of someone's chest and put it in a sterile dish; it may still function (work, operate) normally for a time, but no blood is passing through it. Put a car up on a lift for inspection, and it may function just fine, even if it is not hurtling you down the highway at 65 miles per hour. Fulfilling a function is one thing, functioning normally is another, despite their explanatory relation.

Different items fulfil their functions at different rates, even in normal conditions when functioning normally. The function of sperm—or the sperm-producing device—is to fertilize eggs.<sup>11</sup> But sperm hardly ever do that, even in normal conditions, functioning as normally as can be. Sperm need only fertilize eggs *often enough*. And often enough, for sperm, is not very often. But at the other end of the spectrum some devices fulfil their functions all of the time when functioning normally in normal conditions. The heart provides a good illustration; it pumps blood all the time. Less reliable hearts were selected out. Some items thus have *reliability* vis-à-vis a certain effect as their function. The heart clearly does; the heart is for *reliably* pumping blood.

Functional items often have a plurality of functions. One item may be intended, designed, or selected to do many things. It may have more than one causal-role capacity. Most do. It is thus misleading to speak of ‘the’ function of an item. When it comes to etiological functions and biological items in particular,

156

there is often a plurality at ‘horizontal’ levels on the ‘vertical’ dimension. An example will show what I mean. The tongue is for eating; it moves food around as we chew, and also assists in swallowing. But the tongue is also for talking; it helps form word sounds in conjunction with other parts of our mouth and vocal tract. These are all functions of the tongue at the same ‘horizontal’ level. By doing all of these things, the tongue assists in fulfilling other functions, such as getting enough food and nutrients into the bloodstream so as to nourish other parts of the body, so as to remain healthy and fit, so as to assist in dealing with the creature's natural habitat, so as to find more food, avoid predators, find mates, and so on, so as to live long enough to survive to reproduce. All of these other more distal effects of the tongue are among its functions. On this ‘vertical’ dimension, the tongue has many functions. (And so my title is a bit misleading. I argue for *a* function of comprehension and not for its *one and only* function.)

Items also lose e-functions over time (the appendix), acquire additional e-functions over time (feathers), and even lose one e-function only to acquire another (limbs). When we speak of the e-functions of an item, we often mean current functions (Millikan 1984, 2002; Godfrey-Smith 1994; Schwartz 2002). Origins are one thing; functions (sometimes) are another.

To sum up, there are (at least) four kinds of functions: designed, intention-determined, causal-role, and etiological functions. They all differ notionally and extensionally, though in particular cases two or more may overlap. When it comes to items with etiological functions, *ceteris paribus* they fulfil their functions in normal conditions when functioning normally. Normal functioning is thus operating or working in a manner that, in normal conditions, leads to function fulfilment. If you know an etiological function of an item, you can thereby partially individuate and explain normal functioning in terms of that function. And some items—like the heart—have reliability vis-à-vis an effect as a function. Normal functioning, for the heart, is thus individuated and understood, in part, in terms of reliably pumping blood.

### 3. EPISTEMIC ENTITLEMENT

I now use the etiological notions of function and normal functioning to explicate (a kind of) epistemic entitlement (cf. note 3). I argue that entitlement attaches to beliefs in virtue of the normal functioning of the belief-forming process when the process has forming true beliefs reliably as an etiological function.

First, positive epistemic standings are goods, successes, fulfilments, or achievements understood in terms of promoting true belief and avoiding error. Warrants are positive epistemic standings that epistemically support or ground beliefs. Warrants are positive epistemic standings, and so are themselves goods, successes, fulfilments, or achievements understood, at least in part, in terms

157

of promoting true belief and avoiding error. Any explication of a kind of warrant must explain why the warrant is a success, good, fulfilment, or achievement understood in terms of promoting true belief and avoiding error.

Second, function fulfilment is a success or good. The function sets a standard or norm. Fulfilling the function achieves that standard. Meeting the standard is success or fulfilment, and as such is a good for the item with the function. And failure to fulfil the norm is just that, failure. Normal functioning is also a fulfilment. Functioning normally is functioning (operating, working) the way the item is supposed to operate. Think of the heart. Beating normally is the way the heart is supposed to beat. In beating that way, it pumps blood. Beating normally fulfils a standard for the heart. And failure to operate normally is just that, a failure. Normal functioning is a success or fulfilment for the functional item, and as such is a good for the item. Using ‘achievement’ broadly, function fulfilment and normal functioning are both achievements for functional items.

Despite their explanatory interrelation, function fulfilment and normal functioning are *distinct* achievements. Recall that you can fulfil your function without functioning normally, and vice versa. A heart in a sterile dish may function (work, operate) normally without blood passing through. A car up the lift may function just fine without taking you where you want to go.

Third, for any belief-forming system or process whose etiological function is to form true beliefs reliably, *ceteris paribus* it will form true beliefs reliably in normal conditions when functioning normally. This follows automatically from the logic of etiological functions. Normal functioning for such a process is thus functioning in a way that, in normal conditions, reliably produces true beliefs. Normal functioning for any such system is then individuated and understood in terms of reliability and truth. Again, all of this follows from the logic of e-functions, as just explained.

Putting two and two together, we arrive at an account of epistemic entitlement. Entitlement is a kind of warrant that does not depend on reasons or evidence available to the subject. Entitlements are psychologically and developmentally more primitive kinds of warrants. *Warrants* generally are *fulfilments, successes, or goods* understood, at least in part, in terms of promoting *true* belief and avoiding error that ground or support belief. *Normal functioning* is an *achievement, success, or good* understood in terms of *function fulfilment*. So for any belief-forming process whose e-function is *reliably producing true beliefs*, normal functioning for that process is an achievement understood, at least in part, in terms of promoting true belief and avoiding error. When the belief-forming process functions normally, it epistemically grounds or supports the belief it thereby causes or sustains. A kind of warrant thus attaches to belief when formed on the basis of a normally functioning belief-forming process whose function is to form true beliefs reliably. Since this warrant does not depend on reasons or evidence

158

available to the subject, it is (a kind of) entitlement. A normally functioning belief-forming process that has forming true beliefs reliably as a function confers entitlement on the beliefs it normally causes and sustains. I have established my first thesis, an account of (a kind of) entitlement.

What does this account imply about the massively deceived brain-in-a-vat? In such a case, the standard belief-forming processes are all functioning normally, but they are not reliable. Envatted, the subject's belief-forming process may all function (operate, work) normally, even while not in normal conditions. They thus enjoy (a kind of) entitlement just as if the subject were not envatted. The brain is like a car up the lift for inspection. The subject's belief-forming processes are operating just fine, even though they are not producing and sustaining true beliefs at, so to speak, sixty-five miles an hour.

And what does this account imply about the beneficiary of an accidentally acquired, merely reliable belief-forming process, a process the subject knows absolutely nothing about (that he has it, that it is reliable, and so on)? Merely accidentally acquired, the

process has no function. It's *just* a mutation. Nothing then counts as normal functioning for the process. Though the subject's beliefs so caused enjoy the epistemic good of being mostly true, they lack the achievement conferred by function fulfilment, for it lacks a function, and they lack the achievement conferred from normal functioning, for nothing counts as normal functioning.

Space does not allow me the luxury of responding to familiar objections to etiological accounts of epistemological properties. I hope I have said enough here to make my account of entitlement *prima facie* plausible. Elsewhere I answer some of the more prominent objections.<sup>12</sup>

Perceptual beliefs paradigmatically enjoy entitlement, for (most) perceptual belief-forming systems have reliably forming true beliefs as an etiological function. I argue that comprehension-with-filtering-based beliefs enjoy entitlement for the very same reason; comprehension with filtering has forming true beliefs reliably as an e-function. When our capacity to comprehend and filter functions normally, it meets a standard understood in terms of promoting true belief and avoiding error. That is why our comprehension-with-filtering based beliefs enjoy *prima facie pro tanto* entitlement. (As I noted, this a good place to take a break.)

#### 4. AN ETIOLOGICAL FUNCTION OF ASSERTION

But before arguing that *comprehension-with-filtering* has *reliably* causing and sustaining true beliefs as an e-function, I first argue that *assertion* has inducing (and sustaining) *true* hearer belief as an etiological function. (Recall I am using 'assertion' broadly.) I take up assertion first for two reasons. Making the

159

case for assertion practically makes the case for comprehension-with-filtering; if the argument is plausible for assertion, it will be even more plausible for comprehension-with-filtering. And most of the relevant discussion in the literature so far has focused on assertion instead of comprehension.<sup>13</sup> As a result, the question whether *comprehension-with-filtering* is for *reliably* forming true beliefs is apt to be confused with whether *assertion* has that function. Discussing assertion first should head off potential misunderstanding.

To show that assertion has inducing true hearer beliefs as one of its etiological functions, I begin with an argument from C. A. J. Coady intended to demonstrate the reliability of any practice of reporting. Though Coady's argument unsurprisingly falls short, it reveals an important fact about assertion.

Coady conducted a thought experiment involving a community of Martians. 'Let us suppose for the moment,' he writes,

that they have a language we can translate (there are difficulties in this supposition . . . ) with names for distinguishable things in their environment and suitable predicative equipment. We find however, to our astonishment, that whenever they construct

sentences addressed to each other in the absence (from their vicinity) of the things designated by the names, but when they are, as we should think, in a position to report, they seem to say what we (more synoptically placed) can observe to be false. (1992: 85)

Coady then asked whether any hearer would have any reason to rely on what a speaker says in such a situation. He replies that

any Martian has four powerful reasons for not relying on what others appear to be telling him: (i) he finds their 'reports' false whenever he checks personally on them, (ii) he finds reliance upon them consistently leads him astray in practice, (iii) he finds himself utterly unreliable in what he tells others and it is, at least in part, possible that he is not atypical, (iv) others often give chaotically different reports on those matters beyond his checking.

(1992: 87)

Coady inferred there would then be no practice of making and accepting reports. Given what each Martian can expect of the other,

it is . . . very hard to imagine the activity of reporting in anything like its usual setting with the Martians, for there would surely be no reliance upon the 'reportive' utterances of others. [T]he Martain community cannot reasonably be held to have the practice of reporting . . . (1992: 87)

Coady concluded that any practice of reporting must be reliable.

I think Coady missed a premise. With the premise he needs, Coady's argument goes like this:

(1) If everyone in a community, when making reports, always makes false reports, then no one will have any reason for relying upon and accepting those reports. They have 'four powerful reasons not to do so'.

160

(2) A central reason (motive) for making reports is to get others to believe them.

(3) But if no one accepts those reports, then speakers will have no reason to think that making those reports will lead hearers to believe them. (This is the missing premise.)

(4) So if no one accepts reports, speakers will stop making them.

In short: if all falsehoods uttered, then no acceptances; if no acceptances, then no reason to report; if no reason to report, no reporting. Contraposed: if reports, then not all reports are false. Hence: if reports, then truths.

In 'The Reliability of Testimony' (Graham 2000c) I showed that this argument falls short of establishing the *reliability* of *any* practice of reporting. At best it shows that reporting practices exist because not all of the reports made are false. And 'not all are false' surely does not entail 'most are true'.

Though this argument does not show what Coady hoped, it does bring out a very important fact about linguistic practices: speakers and hearers both need some reason



(motive) to participate. Speakers, presumably, benefit in some way by affecting hearers. If hearers receive no benefit from being so affected, they will probably stop responding in the desired way. So unless hearers get something out of accepting reports, they will not accept them. And if they will not accept them, speakers will not benefit by making them. Then they will not get made. Hearer benefits (partly) explain speaker production.

Ruth Millikan exploits this very fact when arguing that various linguistic devices have etiological functions. She has argued in general that (i) language devices produce effects that interest speakers often enough to encourage continued replication by those speakers *only if* hearers produce hoped-for responses often enough, and (ii) that hearers will continue to produce those hoped-for responses often enough *only if* the results are of interest to hearers (Millikan 2004: 25). Language devices exist, as a matter of fact, only if their use benefits both speakers and hearers.

Applied to (declarative sentences in) the indicative mood, the device we paradigmatically use to make assertions, she has argued that one of its central functions is to produce true beliefs in the propositional content(s) of the utterance, for it is the effect that, often enough, benefits hearers. Truth about relevant subject matters clearly benefits hearers. Dan Sperber says that:

From the point of view of receivers, communication, and testimony in particular, is beneficial only to the extent that it is a source of genuine (and of course relevant) information. Just as in the case of individual cognition, there may be cases where biases in communicated information are beneficial (think of exaggerated encouragement or warnings, for instance), but these cases are marginal. (2001: 404)

Though inducing false beliefs sometimes benefits speakers, and it may also sometimes benefit hearers, if hearers never got true information from assertive utterances of declaratives in the indicative mood, they would quit forming beliefs

161

in response. And then speakers could not even get across false information when they wanted to. (This is what Coady's thought experiment brings to the fore.) And if hearers quit forming beliefs in response, then speakers would quit making those utterances. There would be no point. True hearer belief thus explains why the indicative mood persists. True hearer belief is thus a function of the indicative mood:

If no true beliefs ever resulted from hearer interpretations of indicative sentences, it is clear that indicative syntactic patterns would cease to be used first by hearers and, as a result, by speakers in the ways they now are. . . . The focused stabilizing function of the indicative mood is thus the production of a true hearer belief. 'The function' of the indicative mood is to convey information. (Millikan 1984: 58–9)<sup>14</sup>

Millikan calls this function a stabilizing one, for it 'encourages speakers to keep using the device and hearers to keep responding to it with the same (with a stable) response' (2005: 94).<sup>15</sup>

Since the default, standard use of a declarative sentence in the indicative mood is to assert, and since the same reasoning Millikan offers applies to assertive communication generally, I conclude that an *etiological* function of assertion is to induce true belief in hearers. Truth is a price speakers pay that keeps assertion in play. Though Coady failed to establish the necessary reliability of any practice of reporting, he unearthed an important fact that helps show that inducing true belief in hearers is an etiological function of assertion.

## 5. ALTERNATIVE FUNCTIONS

Some will resist; assertions may regularly *cause* true beliefs—assertion may even as a matter of fact cause *mostly* true beliefs—nevertheless inducing true belief is *not* what assertion is *for*. In this section I reply to four counter-arguments. Responding to these should help persuade the unpersuaded. The last counter-argument will also take us from the function of assertion to the function of comprehension-with-filtering.

162

The first counter-argument focuses on what the speaker has in mind when making her assertion. Call this the *argument from occasion specific intentions*. The argument goes as follows. First, a function is a goal or purpose. Second, the function of a particular assertion on a particular occasion is the goal, purpose, or intention of the speaker on that occasion. Third, speakers have all sorts of purposes on various occasions. Sometimes they intend that their interlocutors form true beliefs, sometimes they only intend shared belief, sometimes they don't intend belief at all, sometimes they intend false belief, and sometimes the truth status is just irrelevant. True hearer belief is thus not the intended effect on each and every occasion. Hence true hearer belief cannot be what assertion, as a type, is for. Assertion is for whatever speakers use it for.

Though I think this argument may be in the mind of many readers, it is rather easy to rebut. From the fact that a particular speaker does not fail at fulfilling *his* aim when making a false assertion, it does not follow that the *etiological* function of the linguistic device he used performed *its* function.

Consider the heart. Though its function is to pump blood, a malicious surgeon may use one as a paperweight. Or consider a bird's wing. Its e-function is to enable flight, though you can cut it off and use it as a fan. Cut off from the bird, it clearly does not perform its function, even if it keeps you cool. A speaker 'can use a language device, just as he can use a tool or one of his body parts, to perform a function or to serve a purpose that is not, however, the language device's *own* function' (Millikan 1984: 52). E-functions and i-functions are simply different things. So from the fact that a speaker does not intend to convey truth on a particular occasion when asserting, nothing follows about the e-functions of assertion.

There is, however, a connection between speaker purposes and the function of assertion. Speaker purposes must be satisfied often enough, or they will stop making assertions. So, in a clear sense, the fulfilment of speaker intentions in general is an etiological function

of assertion. But the Coady–Millikan argument places a constraint on speaker purposes. Speakers must purpose true belief in hearers *often enough* for true belief to be a function of assertion. Here again is Millikan:

But although the stabilizing function of a language device is independent of the purpose of the particular speaker who utters it, it is not independent of speaker purposes in general. The survival of a public language element without change of function must depend upon their being a critical mass of occasions upon which speakers and hearers use the element such that it performs its stabilizing function. (Millikan 1984: 53)

And so long as speakers purpose true belief *often enough*, hearers will respond to assertions with belief *often enough* for speakers to keep making assertions. So speakers satisfy their purposes in general when making assertions because they purpose true hearer belief often enough. Speaker benefits are thus intertwined with hearer benefits, the benefit of true belief in particular.

163

The next counter-argument exploits the standard Gricean view of communication. Call it the *argument from successful communication*. On the Gricean view, communication succeeds or occurs when the hearer recognizes the speaker's communicative intention. A communicative intention is the speaker's intention that the hearer recognize that the speaker intends that the hearer form a belief, answer a question, follow a command, and so on. When the hearer recognizes the communicative intention communication occurs, even if the hearer does not form the belief, answer the question, or follow the command.

Here is Stephen Levinson from his classic textbook:

[C]ommunication involves the notions of intention and agency, and only those inferences that are openly intended to be conveyed can properly said to be communicated. . . . Communication consists of the 'sender' intending to cause the 'receiver' to think or do something, just by getting the 'receiver' to recognize that the 'sender' is trying to cause that thought or action. So communication is a complex kind of intention that is achieved or satisfied just by being recognized. [ . . . ] Attaining this state of mutual knowledge is to have successfully communicated. (Levinson 1983: 15–16)

Kent Bach puts the point more succinctly:

Communicative success is achieved if the speaker chooses his words in such a way that the hearer will, under the circumstances of utterance, recognize his communicative intention. ( . . . ) An act of communication is successful if whoever it is directed to recognizes the intention with which it is performed. In short, its fulfilment consists in its recognition.

(Bach 2006: 153)

In a simple case, in assertively uttering U, the speaker intends that the hearer recognize the speaker's intention that the hearer believes the propositional content(s) of U, and perhaps various implicatures. Communication 'succeeds' if the hearer recognizes the

intention. A speaker can thus communicate when asserting *P* whether or not *P* is true, and whether or not his interlocutor comes to believe *P*. That is simply how communication works.<sup>16</sup>

With this view of 'communicative success' in hand, we can state two related arguments for the conclusion that the function of assertion isn't inducing true hearer belief. The first focuses on the transpersonal communicative process. The second focuses on the speaker's side of the equation. I'll state and discuss the first before turning to the second.

So first think of communication as a transpersonal process: if the hearer correctly recognizes the speaker's communicative intention, then communication has occurred. If uptake did not occur, neither did communication.

164

Here is the counter-argument. First, function fulfilment is a success; failure to fulfil a function is just that, failure. So we can derive the function of communication from what counts as communicative success and failure. Third, communicative success involves recognition of the speaker's communicative intention, and does not require forming belief in the content of the assertion, let alone forming true belief. So if communication can succeed (occur) without conveying true belief, true belief cannot be what it is for; inducing true hearer belief cannot be a function of assertion.

This argument is a non-starter. For it conflates the *ontological* question of what communication *is* (what it consists in, what must happen for it to occur) with the *teleological* question of what it is *for*. The relevant sense of 'success' and 'failure' in the third premise is ontological or constitutive. But the relevant sense of 'success' and 'failure' in the first premise is teleological or purposive, where functions may go unfulfilled. To say what assertive communication is *for* requires pointing to some effect of successful communication. This argument is silent on that issue. Nothing in the argument shows that inducing true beliefs is not what communication is for.

This problem, however, may be avoided if we formulate the argument in the second way mentioned, in terms of the speaker's communicative act itself and its function. Here we ask what the function of the speaker's communicative act happens to be, an act that may fail in fulfilling its aim. The counter-argument would then go as follows. First, function fulfilment is a success; failure to fulfil a function is just that, failure. So we can derive the function of an assertive communicative act from what counts as success and what counts as failure. Third, the speaker's assertive communicative act—his assertion—succeeds when the hearer recognizes the speaker's communicative intention. The etiological function of assertion would then be to induce hearer recognition of the speaker's communicative intention. And recognizing this intention does not require that the hearer believe what the speaker asserted, let alone require that the hearer believe something true. So inducing true hearer belief cannot be what assertion is for. Paraphrasing Paul Faulkner, *when a hearer fails to form a true belief from communication, there is no sense in which the speaker has failed to communicate* (2000: 587).

Though this version of the argument from communicative success focuses on assertion and its effects, and thereby avoids the rejoinder just made to the first version of the argument, it fares no better. Simply put, it doesn't focus on effects that explain why assertion as a type persists. An etiological function is an effect of a type that occurs often enough that explains why the type persists. This argument does not address this issue. All it points out is that a speaker, when making an assertion, intends that the hearer recognize his communicative intention, and when the hearer does, the speaker has communicated. But what is all of this for? What effect explains why speakers make assertions? What effect explains why speakers publicize communicative intentions? What effect explains

165

why hearers even bother to recognize communicative intentions? The argument, though arguably correct about the mechanics of assertive communication, does not address the etiological question of what assertion is for.

Levinson himself is clearly aware of the fact that an *account* of how communication works is one thing, and what linguistic devices are *for* is another. Well after arguing for a Gricean account of communication, he notes that

[since] nearly all of the world's languages have the three basic sentence types: imperative, interrogative and declarative . . . one might argue . . . on the grounds that these seem to be used paradigmatically for ordering, questioning, and asserting . . . that they recur in the languages of the world because humans are, perhaps, specifically concerned with three functions of language in particular—the organization of other persons' actions, the eliciting of information, and the conveying of information. (1983: 40)

The mechanism of communication is one thing; its purpose is another. Griceans may be right about how communication works, while Darwinians (so to speak) are right about what communication is for.<sup>17</sup>

Another counter-argument comes from a paper by Gloria Origgi and Dan Sperber entitled 'Evolution, Communication, and the Proper Function of Language' (2000). They target Millikan's view that the indicative mood is for inducing true hearer belief. They argue that its function is to provide evidence to assist the hearer in recognizing the speaker's communicative intention. The 'function of linguistic utterances', and so the function of assertions in particular, is not to affect hearer beliefs in the content of what is asserted, but rather to provide 'highly precise and informative evidence of the communicator's intention' (Origgi and Sperber 2000: 163).

The following passage contains Origgi and Sperber's reasoning. I have added emphasis to draw attention to a key premise:

Our disagreement with Millikan has to do with the level of processing at which linguistic devices elicit *the reliable response to be identified as their direct proper function*. For Millikan, this *reliable response* is to be found at the level of belief or desire formation, or even at the behavioral level in the case of compliance . . . What is the alternative?  
Linguistic

devices produce *highly reliable responses*, not at the level of the cognitive outputs of comprehension such as belief or desire formation, and even less at the level of behavioral outputs such as compliance, but at an intermediate level in the process of comprehension. Linguistic comprehension involves, at an intermediate and largely unconscious level, the decoding of linguistic stimuli that are then used as evidence by the hearer, together with the context, to arrive inferentially at the speaker's meaning . . . Linguistic devices have proliferated and stabilized because they cause these *highly reliable cognitive responses* at this intermediate level. Linguistic devices provide speakers and hearers with informationally rich, highly structured, and reliably decoded evidence of speaker's meaning.

(Origg and Sperber 2000: 163)

Call this the *argument from highly reliable effects*. The argument goes like this. First, etiological functions are highly reliable effects. Second, the highly reliable effect of assertions (and other linguistic devices and speech acts) is not true hearer belief but the providing to hearers of evidence for the speaker's communicative intentions. Hence providing evidence of the speaker's communicative intention, or perhaps the hearer's decoding and recognizing that evidence (Origg and Sperber speculate on where to locate 'the function' in the causal chain) is the etiological function of linguistic devices.

The argument makes two mistakes. First, as we have seen, etiological functions are not necessarily highly reliable effects. True, pumping blood is a highly reliable effect of the heart, and the function of the heart is to pump blood. But this is not true of all items with etiological functions. Recall sperm. Millikan would be the last person to identify functions with highly reliable effects. Though some items do have reliability as an aspect of their function (like the heart), some clearly do not. You thus cannot infer function from highly reliable effect. Indeed, some highly reliable effects are not even functions. The heart makes noise just as frequently as it pumps blood, and making noise is not one of its functions.

Second, even if providing evidence of the speaker's communicative intention is *an* etiological function of a linguistic device, it does not follow that it is the *only* etiological function. Recall that items with etiological functions often have a number of functions, both horizontally and vertically. Just as the heart assists in removing wastes and circulating nutrients by pumping blood, and it pumps blood by beating at a certain rate in a certain way, so too assertion may induce true hearer belief by providing evidence of speaker's intentions. Assertion, like any other functional item, can have more than one function. Linguistic devices may have proliferated and stabilized (in part) because they provide evidence of speaker's intentions, but they have also proliferated and stabilized by having effects on hearers that explain why hearers care to recover speaker's intentions from the evidence provided in the first place. Origg and Sperber have provided no reason for thinking that inducing true hearer belief cannot be *an* etiological function of assertion. At best they are correct that inducing true hearer belief is not its *only* function. And so

this argument too fails to undermine the thesis that inducing true hearer belief is an etiological function of assertion.

167

Though I have already expanded at some length defending this thesis, I would be amiss to not address one last counter-argument. The argument derives from discussions of the evolutionary origin of the human language faculty. It is widely held in those discussions that language arose in order to facilitate social bonding and social interaction, so that our primate evolutionary ancestors might operate effectively in larger groups, where larger groups promote higher levels of fitness. Robin Dunbar argued for this view at length in his book *Grooming, Gossip, and the Evolution of Language* (1996). Here is Robbins Burling stating the older view of the origins of language in his book *The Talking Ape*:

There was a time, not so many decades ago, when anyone who thought at all about the evolution of language took the reasons for its selection to be more or less self-evident. We use language today for every sort of practical purpose, to exchange information, to coordinate our daily lives, and to instruct the young, so it only seemed reasonable to suppose that, from its earliest days, hunters and gathers used language for much the same purposes. (Burling 2005: 181)

He goes on to embrace the ‘social interaction’ view:

I share a growing consensus that our intelligence evolved primarily as a means for dealing with other individuals. Ours is primarily a social intelligence, and language is one product, and one part of, that intelligence. In preurban societies, and even our own, language in its most delicately nuanced form is used, not so much for basic subsistence tasks, as for establishing, maintaining, and refining social relationships. It is when dealing with people, not with material objects, that we call on our richest linguistic resources.

(Burling 2005: 184)

And so we once thought that language arose for ‘exchanging information’ but we now think language arose for ‘dealing with other individuals’. Communication, and assertion in particular, is thus not for inducing true hearer belief, but rather for interacting with others and maintaining social bonds. So instead of focusing on more proximal effects of assertion (such as providing evidence of speaker intentions or the actual recognition of such intentions), this argument focuses on the more distal effects of inducing belief through communication. Assertion is thus for social bonding, not for conveying true information. Call this the *argument from social interaction*.

The error in this argument should be easy to spot. There is no opposition between social bonding and conveying information as etiological functions for assertion. A closer reading of Dunbar and Burling (see also Pinker and Bloom 1990; Dessalles 2007; Hurford 2007) reveals no such opposition at all. In fact, Dunbar holds that language conduced bonding and so the forming of larger groups *by* conveying information about social matters. Social bonding is simply higher up in the vertical hierarchy of language functions. Conveying true information is one etiological function of communication:

language evolved to facilitate the bonding of social groups, and . . . it mainly achieves this aim by . . . the exchange of socially relevant information. (Dunbar 1996: 120)

168

language evolved to exchange information about other individuals in the group . . . language evolved to exchange social information. (Dunbar 1996: 142)

Language . . . is very much a social tool . . . it allows us to exchange information relevant to our ability to survive in a complex, constantly changing social world . . .

(Dunbar 1996: 170)

And so the old view is that language evolved primarily to convey information about where and how to find food and avoid danger. The new view is that it evolved in addition, if not primarily, to convey information about social matters to build groups where building groups, in turn, has obvious implications for survival. Burling concludes:

Thanks to the eagerness of people to display their knowledge, we are able to learn important things from conversation and it is people that we learn most about. Since it is maneuvering through the social system that gives us the most difficult challenges of our lives, the information we gain by learning how others have behaved is extremely valuable. It is fortunate for the hearer, then, that we are at least as eager to pass on information as to receive it. (Burling 2005: 196)

Either way, language originated, in part, to convey true information.

Summing up, assertive utterances provide evidence of speaker's intentions. Hearers recognize speaker's intentions and thereby comprehend speaker's assertions. Comprehension in turn induces true belief, often enough, and thereby provides an obvious benefit to hearers. Conveying true belief builds social bonds, where building bonds conveys benefits to both speakers and hearers. Conveying true information is one effect of assertion that clearly explains why assertion, as a type, persists. 'The immediate effect of language when used in standard conditions, in other words, its proximal function, is . . . to convey . . . the information and knowledge possessed . . . by the speaker' (Dessalles 2007: 315).<sup>18</sup> (In need of a break? This would be another good place to pause.)

## 6. COMPREHENSION, RELIABILITY, AND FILTERING

Absent counter-considerations, we take comprehension at face value. My topic is the epistemology of this process. What kinds of warrant do beliefs

169

based on comprehension-with-filtering enjoy? I argued a priori in section 3 that beliefs based on a normally functioning process that has forming true beliefs reliably as an etiological function enjoy *prima facie pro tanto* entitlement. Beliefs based on comprehension-with-filtering will thus enjoy entitlement provided comprehension-with-



filtering has forming true beliefs reliably as an etiological function. Does comprehension-with-filtering have that function?

I have just argued that the usual causal precursor to comprehension—assertion—has inducing true hearer belief as an etiological function. If assertion is for producing true belief, and it does so by triggering comprehension states, so comprehension states are also for inducing true hearer belief. If not enough of the beliefs formed on the basis of comprehension were true, then hearers would not be moved to comprehend assertions or transition from comprehension to belief in the first place. The argument of section 4 and its defence in section 5 carries over automatically to comprehension. And since the transition from comprehension to belief is not intentional in the way assertions are, the first three counter-arguments just discussed do not touch the extension of the argument from assertion to comprehension. Comprehension has inducing *true* beliefs as a function.

But what about inducing true beliefs *reliably*? Does comprehension have producing true beliefs *reliably* as a function?

I am happy to grant the possibility that comprehension taken alone, comprehension without filtering, may fail to have inducing true beliefs reliably as a function. I am happy to grant the possibility for my thesis does not require it. My thesis is that comprehension-*with-filtering* has forming true beliefs reliably as a function, *not* comprehension *neat*, comprehension *taken alone*. It's the filtering, or so I argue, that is for producing a sufficiently high truth ratio. On the other hand, if it turns out that comprehension neat reliably induces true beliefs, and that is why we do it, then comprehension as such would confer entitlement when functioning normally. My argument here would then establish a higher degree of entitlement for beliefs based on comprehension-with-filtering.

But before I argue that comprehension-with-filtering is for reliably inducing true belief, let me make one point very clear. I am not arguing a priori that comprehension-with-filtering is necessarily reliable. My argument is broadly empirical. What's a priori is my account of entitlement. Its application to comprehension-with-filtering is empirical (cf. Goldman 1986).

So why should we think comprehension-with-filtering induces true beliefs *reliably*, and why should we think it has inducing true beliefs reliably as an etiological *function*? The answer turns on understanding why we filter, what filtering is for.

Recall the logic of the argument derived from Coady and Millikan. Speakers assert for some benefit. Hearers in turn form beliefs for some benefit. Unless hearers formed true beliefs often enough, they would stop forming beliefs as a response to speaker assertions. Speakers would no longer benefit by making assertions. Assertions would then no longer get made. Inducing true hearer belief is thus a stabilizing function of assertion. So far I have emphasized how hearers benefit: they form true beliefs. I have not discussed how speakers benefit, benefits that explain why we are so disposed to make assertions in the first place.

There are many ways speakers might benefit by making assertions. By providing useful and relevant information, speakers gain status and build coalitions (Dessalles 2007), build social bonds (Dunbar 1996; Burling 2005), enhance cooperative activities (Tomasello 2008), and even attract reproductive partners (Burling 2005). And the argument so far reveals that speakers acquire these benefits by providing true information, at least often enough. Some of these benefits explain the origins of informative speech and communication. Some explain its current shape. Providing true information leads to speaker benefits that historically and currently explain why assertions get made, and so explains why assertion as a type persists.

So we benefit both as speakers and hearers in a variety of ways from communication, where the passing along of true information plays a vital role in accounting for all of this. But despite the utility of true information, how often do we actually present the truth? How honest are we? Don't we often deceive? What about the actual frequency of cheats and liars? Don't speakers often appear to cooperatively provide useful information that turns out to be bogus? And can't speakers sometimes reap the benefits associated with making assertions but without paying the price of providing true information? Perhaps speaker benefits on occasions may be best achieved not by providing true information, but by distorting the truth, or even saying what is clearly false (Sperber 2001). And won't the cost to speakers of acquiring true and useful information, and the cost of sharing such information, sometimes serve as a motive for trying to reap the benefits of being a good information provider by providing apparently genuine information that is not the real thing (Dessalles 2007)? Won't conflicting motives and the cost of acquiring information lead speakers, on occasions, to say what is not the case? Isn't some degree of misinformation guaranteed by the logic of the argument? Dan Sperber concludes that

a significant proportion of socially acquired beliefs are likely to be false beliefs, and this not just a result of the malfunctioning, but also of the proper [normal] functioning of social communication. . . . [T]he cognitive manipulation of others is one of the effects that makes the practices of testimony and argumentation adaptive. This contributes to explaining why these practices have evolved and stabilized among humans.

(Sperber 2001: 402)

Some degree of misinformation is thus bound to result from the very logic of the argument that shows that assertion has providing true information as (one of) its etiological function(s).

On the other hand, hearers do not sit idly by as speakers mislead or otherwise misinform. Hearers develop countermeasures. Hearers develop capacities for filtering out reliance on fake information providers. Indeed, for communication to stabilize, hearers must develop such filters. Here again is Sperber:

if communication has stabilized among humans, it must be that there are ways to calibrate one's confidence in communicated information so as that the expected benefits are greater than the expected costs. (Sperber 2001: 406)

One important filter consists in detecting cognitive conflicts: coherence checking. Sperber suggests that

coherence checking—which involves metarepresentational attention to logical and evidential relationships between representations—evolved as a means of reaping the benefits of communication while limiting its costs. (Sperber 2001: 410)

Jean-Louis Dessalles agrees with the challenge and its solution:

Let us assume that whatever benefit speakers derive will increase with the salience of the situations they speak of. This makes for a strong temptation to exaggerate or even to tell lies. The risk run by a hearer is that of affording an undue benefit to the speaker. Against this risk, hearers have two strategies, trivialization through comparison with already known situations and detection of inconsistencies. Both of these rely on faculties with which natural selection has endowed us as human beings, among which are a sense of probabilities and the ability to detect cognitive conflicts. Each of these conversational mechanisms has thus much the same function as the other, that is avoidance of cheating in communication. (Dessalles 2007: 331)

hearers . . . run the risk of being cheated in verbal interactions . . . The original role of the detection of cognitive conflicts can be made sense of as the attempts of hearers to guard against lying. (Dessalles 2007: 330)

Though there are incentives for speakers to cheat or mislead, we develop filters for screening out their influence. Coherence checking is one. We encountered other filters in the first section. We deploy all of these and more.<sup>19</sup>

172

Filtering thus plays a very important role; it ensures the reliability of beliefs formed on the basis of comprehension, for misleading or untrustworthy assertions often get filtered out. We're not so easy to mislead:

hearers have effective ways of assessing the quality of information supplied to them and . . . this makes it difficult to mislead them. (Dessalles 2007: 332)

the only signals that natural selection can favour are the reliable ones. With language . . . the reliability of the message is . . . guaranteed by . . . its resistance to hearer's assessments. . . . It is not . . . easy to tell lies, even with words, as hearers test the logical consistency of what they are told. (Dessalles 2007: 331)

Furthermore, filtering not only dampens the possibility of accepting a false report, it also provides an incentive for speakers not to cheat in the first place. For cheaters often get caught. Indeed, a good deal of the social gossip hypothesis is driven by the idea that one very useful kind of information we receive from communication is information about

which speakers tend to cheat (Dunbar 1996). Given the cost of getting caught, nature has even implanted in speakers a disposition to stick to the facts:

[We have] an emotional warning whenever we are tempted to lie. Our own nervousness cautions us to be careful, and the nervousness may be visible enough to make others suspicious. We are not as good at lying as evolutionary theory and sheer self-interest might lead us to expect, but since lying can be risky, caution may be beneficial. Once people catch you in serious contradictions, and once they start to share this interesting and relevant information with your friends, you will lose badly in the competition for prestige. You would have fared even worse in a society of a few hundred people where everyone know and depended on everyone else. Most of us simply don't have the skill to stray very far from the truth without getting caught. We can avoid danger by sticking reasonably close to reality. (Burling 2005: 196)

Filtering not only filters out false or misleading assertions, it also provides an incentive for speakers to not mislead in the first place.

Comprehension-with-filtering may not, as a matter of fact, be as reliable as perception. But for all that, it is still an awfully good way of acquiring information; it is a very reliable guide to the way things are. And the reliability enhancing aspect of filtering explains, at least in part, why we bother to filter what we take others to assert. Comprehension-with-filtering has inducing true beliefs reliably as a function. Getting informed is one thing, among others, that comprehension-with-filtering is for.

There are different kinds of warrant. Many comprehension-based beliefs are supported by justifications. But comprehension-based beliefs also enjoy entitlements. For beliefs based on normally functioning processes that have

173

forming true beliefs reliably as an etiological function enjoy entitlement, and comprehension-with-filtering has that function.<sup>20</sup>

## 9 Knowing from Being Told

Alan Millar

### 1. INTRODUCTION

In this chapter I sketch an epistemology of what I shall call straightforward cases of testimony. Here is an instance. Wondering if I have left a certain book in my study I phone home from the university. My son answers, I tell him about the book, he looks in my study and sees that it's on my desk. He tells me this. From his telling me I come to know that the book is on my desk. Straightforward cases are like this. They have the following features.

- (1) It is entirely natural to suppose that recipients gain knowledge of what they are told from being told it. As a result they can rationally be assured of the matter and cease any enquiry into it that they have been pursuing. Should the need arise they could responsibly vouch for the truth of what they have been told.
- (2) Recipients do not deliberate about the likelihood or otherwise that what they have been told is true. They unhesitatingly accept what they have been told.
- (3) Notwithstanding (2), recipients are not indiscriminating. Features of the utterance and its context are relevant to their unhesitant acceptance of what they have been told and to their knowing what they have been told. In the example I have given it is relevant, for instance, that it is my son who is telling me, that the content of his communication is a routine matter, and that he speaks to me in a straightforward matter-of-fact way. I might not have so readily accepted something he had told me if the subject matter or context had been different. For instance, if he had told me that there is a smudge on my nose, in the context of light-heartedness, I might well have been more circumspect, knowing that he takes pleasure in showing me to be gullible.

The challenge posed by straightforward cases arises from three considerations. The first is that the knowledge gained by recipients is evidence-based if only in the sense that something occurs—an act of telling—that is empirically detected and that supplies the reason or at least part of the reason for accepting the thing told. The second is that telling is a speech act by means of which informants give it to be understood that they are

informing those with whom they are communicating. The third is that recipients need not at any conscious level register the features of the act of telling that account for their unhesitating acceptance of what they are told. Stepping back from the situation in which my son confirmed the location of my book, I can think of factors that help to explain my unhesitating acceptance of what he told me. There are factors relating to the content of the utterance. It was clearly within my son's capabilities to ascertain that the book was in my study. He could have little reason for misleading me on this trivial matter. He speaks as he would if straightforwardly communicating information. The context is one in which I have expressly enquired about the book and would be unlikely to appreciate his fooling around. However, there is no reason to suppose that these factors need have been registered at the level of judgement. Indeed, the factors themselves and their full range are elusive. I cannot describe them specifically but can at best gesture towards them. I could not pin down what it was about his manner of speaking, the context and so forth that made the communication effective. Even to provide a gesturing description requires a degree of reflection that recipients in the straightforward cases are not bound to have.

To some it might sound odd to say that my son's telling me that the book is at home is *evidence* that the book is at home and that the knowledge I gain is, accordingly, evidence-based. Calling it evidence might suggest that it is to be viewed as something to be weighed alongside other considerations for and against taking what I am told to be true. But a feature of the straightforward cases is that no weighing up is done even though knowledge is acquired. Another motivation for pausing over whether knowledge in the straightforward cases should be treated as evidence-based is supplied by the second of the features of those cases noted above. Whereas treating acts of telling as evidence for the thing told might suggest that the epistemology of testimony is reducible to the epistemology of empirical evidence in general, the fact that telling is a speech act suggests that a central role in the epistemology of testimony should be assigned to speech act considerations. If we accede to the latter suggestion then it might seem that we should not treat acts of telling as evidence at all but as calling for special treatment drawing upon the character of the practices surrounding the speech act. Such a reaction would be unjustified. If telling is a distinctive speech act then it might well be that an adequate epistemology of testimony must incorporate elements that have no echo in the epistemology of knowledge based on non-testimonial evidence. But that does not establish that when those

elements are in play we should not think of acts of telling as constituting evidence in the broad sense countenanced above.<sup>1</sup>

A further reason for hesitating over whether knowledge received in the straightforward cases constitutes evidence for the thing told arises from the thought that it would be hard on such a view to account for the recipient's entitlement to treat the act of telling as evidence. The problem that seems to arise here is closely linked to the third of the

features of the straightforward cases that I noted above. This has to do with the elusiveness of the factors that explain the ready acceptance by recipients of what they have been told. There is a way of thinking of knowledge from indicators—indicating phenomena like readings on fuel gauges and tyre marks on roads—on which this elusiveness seems problematic. It can be that the occurrence of tracks of a certain sort on a path indicates, and is in that way evidence, that deer have recently passed. People who know about deer might be able to tell that deer have recently passed from the presence of such tracks on a path. A standard way of thinking about how they can do this is to suppose that they justifiably accept, and exploit, a generalization to the effect that when there are tracks of the relevant sort on a path then deer will have recently passed by. In particular cases they apply the generalization to the circumstances and conclude that deer have recently passed. If we try to adapt this way of thinking to straightforward testimonial cases then we would have it that if you were to accept that  $p$  on the basis of being told that  $p$  by an informant A, you would have to apply a suitable covering generalization to this act of telling.<sup>2</sup> What might the generalization be? Given the elusiveness of the relevant factors it is hard to say. Do we need a generalization about A or about people like A? If the latter, then what similarities are relevant? Do we need a generalization about utterances with the same sort of content? If so, what determines the relevant sort? Analogous problems arise about the sorts of context and manner of speaking that are relevant. And there is the further matter of how we know that all the relevant factors have been covered. The reason why this matters is that on the model under consideration we are supposed to base our acceptance of testimony on facts to the effect that the speaker, context, and so forth have such-and-such features. To the extent that it is unclear which features count for us, and accordingly which generalizations are being exploited, it will be unclear what is the basis of our acceptance. Now, one might take these problems to provide the motivation for further enquiry into the form of relevant generalizations and the features they incorporate. I am sceptical that such a project has much chance of success. This

is in part because I think that a different approach is more plausible and avoids the problems. My aim in this chapter is to outline key features of this account. I shall focus on two in particular. The first is that a central role is assigned to the speech act of telling. I suggest, more specifically, that we must think of telling as a move in a *practice*, where a practice is conceived as an essentially rule-governed activity or cluster of such activities. Crucially, considerations about speech acts do not take us all of the way. This takes us to the second feature—an account of the kind of sensibility that enables us to recognize acts of telling and to recognize some (trustworthy) acts of telling as indicative of the truth.<sup>3</sup>

## 2. SAYING, TELLING, AND THE PRACTICE OF INFORMING THROUGH TELLING

Telling someone that  $p$  is a distinctive communicative act. My telling you that  $p$  is an act of saying to you that  $p$  by which I give you to understand that I am thereby informing you that  $p$ . I take informing you that  $p$  to entail speaking from knowledge that  $p$ , with the aim that you should thereby come to know that  $p$ . Thus I do not inform you that  $p$  unless  $p$ , though I may give you to understand that I am informing you that  $p$  when I know full well that not- $p$ .<sup>4</sup>

Not all cases of saying to someone that  $p$  are cases of telling that person that  $p$ . If by way of advice I say to a colleague that he should consider pursuing some topic, I do not take myself to be informing him of anything and do not expect him to regard what I say as knowledge that I am conveying to him. My aim will be to encourage him to pursue the topic. Good advice recommends a course of action and presents considerations in the light of which the course of action can be seen to be a good idea. In advising my colleague I don't wish him to consider pursuing the topic just on my say so but rather on account of whatever it is that makes acting in the suggested way a good idea.

There are other cases in which sayings are plainly not acts of telling. I might 'sound off' about the qualities of well-known political figures. These may be unqualified and accompanied by manifest strength of conviction. But I would view them as expressions of opinion and perhaps anger too, not as aiming to convey knowledge. That is how others are likely to take them.

I said telling is a move in a *practice*. The practice may be conceived as that of informing through telling, but it should be understood that the practice embraces

178

both informing through telling, understanding acts of telling, and adopting a stance towards what one is being told. In the sense intended here, a practice is *an essentially rule-governed activity or group of interrelated activities*.<sup>5</sup> Playing in a game of football, for instance, is participating in a practice. The activity is subject to the rules of the game and is indeed individuated by those rules.<sup>6</sup> Informing through telling is also subject to rules. I shall assume that the following are rules:

- (A) when informing someone that  $p$  you should speak from knowledge that  $p$ ;<sup>7</sup>
- (B) when informing someone that  $p$  you should intend that this person should, by your act of informing, come to know that  $p$ .

The (A)-rule makes sense of the fact that by telling you are supposed to inform. While an act of telling might not inform, and indeed might be deliberately deceptive, a felicitous act of telling—one in keeping with the rules—does inform. The (B)-rule accounts for a way in which an act of telling can be infelicitous even when the informant speaks from knowledge. Suppose that Bill, speaking from knowledge, tells Sally that her husband John is not having an affair. But Bill's intention is to get Sally to suspect that John is having an affair by trading on Sally's alarm at John's friendship with a female colleague



and her tendency to suspect that he (Bill) would lie to protect his friend.<sup>8</sup> Arguably, Bill's act of telling is infelicitous in that he has no intention to inform Sally. As we might say, that is not what telling is supposed to be; hence the second rule.

How should we think of the relationship of participants in a practice to the rules governing the practice? Consider a game of football again. Obviously, players are subject to the rules in that the rules apply to them. The rules also set limits to what counts as playing the game. Games can proceed despite some rules being broken, but there are limits to how anarchic play can be if a game is to proceed. Whatever players actually do they incur a commitment, just in virtue of being players, to following the rules of the game. The commitments here are normative in that they relate, in a way to be explained, to how a player ought to behave. Some might be attracted by the idea that the rules do not generate any such commitments beyond those incurred by aiming to play. If we ask what reason there is for players to follow the rules, the answer by this account would be, 'because they intend to play the game and could not do so without

179

following the rules at least to some extent'. On this way of thinking the only normativity we need to acknowledge in connection with rule-governed activities is instrumental—following the rules by and large is what you have to do if you are to play. The rules have no normative force; it's just that you won't count as playing unless by and large you conform to them. If this is right there need be nothing wrong with a football player's being in breach of a rule. That is because flouting a rule evidently need not be at odds with being a player and so does not put one in breach of the commitment incurred by intending to play the game—the commitment to doing what is necessary to that end. All we have is a clash between what the player has done and what some rule prescribes. I do not think that this is satisfactory. There is something wrong about being in breach of a rule and this is reflected in the fact that breaches of the rule are subject to legitimate criticism. What is it that makes the criticism apt? Ethics need have nothing to do with it. I suggest that we need the notion that it is in the nature of rules governing a practice that those who participate in the practice incur a (normative) commitment to following the rules. Behaviour that flouts a rule may (legitimately) be criticized because it is in breach of this commitment. That you have incurred the commitment does not entail that you ought to follow the rules. (So by speaking of a *normative* commitment I do not mean to suggest that it imposes an obligation to do that to which you are committed.) There can be games so appalling that no one ought to play them and therefore no one ought to follow their rules. By my account it remains true that if you play such a game you incur a commitment to following the rules. The commitment amounts to something like this: you ought to avoid continuing as a player, and not follow the rules. There are two ways to discharge this commitment: one is to carry it out, that is, follow the rules, the other is to withdraw from the game, thereby removing the condition in virtue of which the commitment was incurred. It could be that you ought to withdraw from the game. On this view the mere existence of a practice gives no one, not even participants, an obligation to follow its rules. The sense that practices do not generate obligations to follow their rules

may explain the resistance that some have to the very idea that there is an intrinsic normative dimension to participation in practices. But we can acknowledge the point about obligation while still acknowledging that practices have an intrinsic normative dimension. What we need is a weaker normative notion than that of obligation—the notion of a commitment.<sup>9</sup>

The notion of a practice is philosophically interesting for at least two reasons. One has to do with the explanation of action. The fact that there is a practice that commits participants to acting in a certain way figures in the explanation

180

of why participants act in that way. This happens in part because people simply become habituated to complying with what the practice requires. The other reason is epistemological. Observers and competent players of a game know how the game is supposed to be played. This enables them to form reasonable expectations about the course of play and to make sense of the play that unfolds. In a team game like soccer, knowledge of what the game requires enables players to coordinate with others on their team and to be prepared for moves on the part of the opposing team. Grasp of a practice—knowing how to engage in it—can ground expectations about how participants are liable to act even in the absence of knowledge of the participants as individuals. This is not to deny that in some games, for instance, tennis, a player's individual quirks and preferences have an important bearing on what in detail he or she will do. The point remains that even then we do not need to know the quirks or preferences to make out the general shape of what a player will do.

The sorts of practices I have mentioned thus far are instituted by design and have more or less well-defined rules by reference to which participants guide what they do. There can be practices for which rules are never formulated but which develop because certain ways of proceeding come to be seen as *the done thing* or *the way to go about things*. This is manifested by the fact that there arise mutual expectations of conformity, and a willingness to do the done thing and to view departures, whether by oneself or others, as open to legitimate criticism. What is it for there to be rules of such a practice, given that the rules are not instituted? Consider, for instance, the practice of giving and receiving invitations. It's a rule of the practice that you don't issue invitations without planning, and implementing plans for, the activity in question. There is a certain pattern of behaviour that is in keeping with the rule: issuing invitations, planning, and implementing plans for, the activity in question. There is behaviour that is not in keeping with the rule: issuing invitations and not planning, or not implementing plans, for the activity in question. To say that there is a rule along these lines is, roughly speaking, to say that behaviours of the former kind are acknowledged by those giving or receiving invitations as right or correct or appropriate, and behaviours of the latter kind are acknowledged to be wrong or incorrect or inappropriate. Part of what these acknowledgements involve is an understanding that one's behaviour is expected to conform to the first pattern and not to conform to the second. These informal practices, as one might call them, have epistemic

utility as much as those that have been instituted by design. That there is a practice of giving and receiving invitations enables those who issue invitations to have a reasonable expectation that those invited will respond in a certain way, for instance, accept the invitation and then turn up at the set time. It enables those invited to have a reasonable expectation that the activity will indeed take place. All this can proceed smoothly with little detailed knowledge of the characteristics of those with whom one is interacting beyond the fact that

181

they may be expected to *do the done thing* with respect to the giving and receiving of invitations.

A very great deal of what we expect of others, and much of our understanding of them, turns on our having a grasp of practices in which we, and they, participate. In our familiar communities, for the most part we are not like anthropologists striving to make sense of what others are about. We have a sense of what makes sense when people are engaged in activities with respect to which there is by custom a right way to go about things. We have such a sense because we, like they, have been initiated into the appropriate ways to go about things, and into thinking of them as the appropriate ways. That initiation brings with it mutual expectations of conformity and preparedness to acknowledge the legitimacy of criticism when there are breaches.

Linguistic practices ensure that there are shared abilities and sensibilities that make communication and mutual understanding possible. The upshot of the shared abilities and sensibilities is that we do not in general approach users of our language with an interpretative problem—what to make of the sounds they utter or the inscriptions they write. We are already geared up to viewing them as fellow speakers of our language and to recognizing what they are doing with the words they use: asking us this, telling us that, and so on.<sup>10</sup> A sense that there are right and wrong ways to communicate with our language serves to reinforce our own continuing conformity to the demands of using our linguistic practices. It is easy to overlook this because so much in our use of language, and our understanding of uses of language, is unreflective. It seems plausible however that our willingness to adjust our usage when we find that it is out of kilter with common practice, and we are not in the business of being innovative, testifies to our having a sense of correctness and incorrectness in linguistic matters.<sup>11</sup> There is a standing temptation to suppose that our conformity to prevailing practices can be explained simply in terms of instrumental reasoning: if we do not conform we will not understand or be understood. My point is that thanks to the existence of practices, and our initiation into them, we do not routinely, and do not have to, get to the stage of reasoning in this way; we are instilled from the start with a sense of there being correct and incorrect ways to speak and write.

Against the background of these ideas I shall now pick up a claim made towards the end of section 1. There I said that my account of the epistemology of testimony would make

use of the idea that our ability to gain knowledge in straightforward cases of testimony depends on our having a certain sensibility. This is the theme that I shall now pursue.

182

### 3. RECOGNITION AND RELIANCE ON INDICATORS

Much of our knowledge is perceptual, being gained directly, without reasoning, through the exercise of perceptual-recognitional abilities.<sup>12</sup> We might have, for instance, the ability to recognize that a sound is that of a fire alarm, or recognize by looking that the flowers in a vase are daffodils, or recognize that a fabric is silk from the way it feels. Possession of any such ability requires more than that the appropriate sensory system is functioning adequately; it requires that we have appropriate concepts *and* have learnt to apply them in response to what we perceive. We have to learn to recognize robins or goldfinches from the way they look. We have to learn to recognize people we know from the way they look or sound or perhaps from their gait. Thanks to perceptual-recognitional abilities, perceptual knowledge can extend far beyond the superficial features of things that go to make up their appearance. We can recognize Bill as Bill and not as just someone or something having the look of Bill. Such abilities involve refined sensibilities by which one is attuned to the significance of a host of physical cues that one would be hard put to describe. This is rather obviously so in the case of face recognition. The information we can capture by a description of a face is a tiny fraction of that to which we respond when we recognize the person whose face it is.

It is thanks to our perceptual-recognitional abilities that we are able to take in 'rich' facts, and not merely facts concerning the superficial features of things. Such abilities enable us to have perceptual knowledge that extends well beyond the features that make up the superficial appearance of things to this or that sense modality. As conceived here the exercise of such abilities *is* the acquisition of knowledge. The ability to recognize a goldfinch from its visual appearance—from the way it looks—is an ability to tell of goldfinches that they are goldfinches from the way they look. I exercise this ability only if I tell, and thus come to know, of the thing looked at, that it is a goldfinch. For present purposes it is especially important to appreciate that the knowledge acquired in virtue of exercising this ability is not implicitly evidence-based but is genuinely recognitional. Though I tell that the bird at which I am looking is a goldfinch from its appearance and might pick out some salient markings, it is the entire *Gestalt* presented by the bird to which I respond and which triggers recognition. Goldfinches have that distinctive flash of red on the side of the beak, but it is (showing you the illustration in the bird book) *that* distinctive flash of red in *that* feathery context. You need to see it (or some representation of it) to know what I am talking about. As I look at a goldfinch I do not propositionally represent its visible features in any very specific way, nor therefore do I work with a generalization that connects

183

specific such features to being a goldfinch. No plausible model of perceptual knowledge of the presence of goldfinches should require this of me.

Perceptual-recognitional abilities have not figured prominently in the epistemology of testimony.<sup>13</sup> That is not surprising. If we gain knowledge that *p* from testimony that *p* it is on the basis of the testimony, which clearly contrasts with perceiving that *p*. Against the background of the traditional problem of other minds, one might suppose that even knowing that we are being told that *p* is inferential—that we judge that we are being told that *p* on the basis of evidence that can be specified independently of the meaning of the words used, and of their being used to tell us something. On this way of thinking, it is assumed that the evidence can comprise only what can be taken in perceptually, and that what can be taken in perceptually is confined to the level of superficial features. It does not include psychologically loaded facts to the effect that someone's words mean this or that, or that he or she is telling us this or that.<sup>14</sup> A problem for such a view is that there is no reason to think that we routinely treat psychology-free facts as evidence for a judgement as to what the speaker is telling us. And if that is so, then such facts can hardly figure routinely in the explanation of how we can know things on their basis. But there is a contrasting picture on which perceptual-recognitional abilities have a role even in relation to knowledge of 'other minds'.<sup>15</sup> For instance, of people we know we can sometimes tell that they are pleased about something from their smile and sparkling eyes. We should resist the temptation to treat these cases on the model of inference from evidence. The cues are too subtle for that. They trigger recognition of a how the person feels rather than furnish evidence for what the person feels. It is more a matter of deploying a sensibility whereby we are responsive to the cues that enable us to recognize what the person is feeling. Focusing on recognitional abilities, here as elsewhere, opens up the prospect of explaining how a host of factors can be relevant to acquiring knowledge that something is so even though the knowledge does not result from reasoning from evidence pertaining to those factors. This is of crucial importance for the epistemology of testimony. If I were to tell you that you left your umbrella in my room, then unless the circumstances were rather unusual, you would thereby come to know that I am so telling you. Your coming to know would be a response to features of my utterance, and of the context, that you would be hard put to articulate. The role of these features, though, is not

that of furnishing you with evidence. Their presence prompts recognition: you would simply hear me to be telling you that the umbrella is in my room. You would not need to base your judgement on evidence any more than you would when recognizing a friend from her appearance or telling that she is pleased from her appearance.

Our concern, of course, is not simply with how we tell when we are being told this or that, but with how, in the straightforward cases, we tell that something is so from being told it. On the picture I am working with someone's telling you that *p* in a case like this

can indicate, and in that way be evidence, that  $p$ . But it being such evidence is obviously only part of the story. The question is what entitles us to treat it as such.

What might be called the standard model of knowledge from indicators (indicating phenomena) envisages that experience furnishes us with facts about particulars, which provide the basis for generalizations, which can then be applied to new cases. On this model, to know that a's being F indicates that it is G one needs empirical support for a covering generalization that connects a's being F with its being G. This might be that everything F is G. Obviously, there can be support for such generalizations, and one can be justified in thinking them true or at least in placing some fair degree of confidence in their being true. I have already pointed to difficulties in applying such a model to testimony but it is open to doubt that it applies to much of our knowledge from non-testimonial indicators. Who among us has tested whether our car's fuel gauge is accurate or whether only cows moo in just the way we think of cows as mooing? Yet we sometimes tell from the reading on the fuel gauge roughly how much fuel is in the tank and we sometimes tell that cows are nearby from the sound of their mooing. Our experience is undoubtedly relevant to our abilities to do such things. I want to suggest that we should think of the role of experience as that of inculcating and shaping the abilities. By 'experience' here I include our perceiving that this or that is so but also the experience by which we are bombarded with observations, corrections, explanations, and so on by those who surround us. Through such experience we get the hang of telling that it has been raining from the wetness of the streets and pavements, or telling that someone wishes to enter from the knocking sound at the door, or telling that it has been very cold from the frost on the grass, and so on. When I say that we get the hang of these things I mean simply that we learn to do them. The upshot of the learning process is a kind recognitional ability. This time though it is an ability to recognize a situation as having a certain significance. We should think of the role of experience accordingly. Its role is in some respects akin to the role of experience in shaping an ability to saw wood straight, or cultivate orchids, or cook fish perfectly. In these latter cases we take steps towards a certain end, and gain feedback from our faltering attempts, perhaps assisted by guidance from those who know about these things. The upshot of the learning process is mastery of a technique. The role of experience is that of honing that technique

185

by trial and error, so that in time it becomes an ability to do something well. The test for whether we have achieved this is whether we saw wood well, cultivate orchids well, cook fish well, as the case may be. Likewise, the test for whether we have mastered the ability to tell that deer have recently passed by along a path from the presence of tracks on that path, is, roughly speaking, whether we pull off this feat on all or nearly all of the occasions on which we aim to do so. Often the abilities are inculcated by others rather than deriving from our own independent enquiry. Sometimes, though, they are acquired through such enquiry. Without guidance from others, a person shipwrecked on an island might learn that the presence of certain droppings indicates that rabbits are around. This will involve observing rabbits produce such droppings and learning to discriminate the

rabbit droppings from others that might be around. Initial steps towards this may be faltering. Perhaps there are droppings very like rabbit droppings, produced by another small mammal. But in time the learner catches on to recognizing just the right sort. The upshot again is a certain technique, in this case judging that rabbits are around on the basis of the droppings. Getting the technique right is acquiring a recognitional ability—the ability to recognize droppings as having a certain significance: that rabbits are around. The test for whether the ability has been acquired is whether over a wide range of cases in which an attempt is made the learner judges correctly with a very high degree of reliability and judges wrongly only when there is no ‘interference’—distraction, inadequate attention, or the like. The focus of the learning process is on identifying the right sort of dropping. The same applies to knocks at the door. Not any sound at the door is the right kind of sound to count as a knock. One has to learn to discriminate the knocks from the bumps and scrapings. Learning to connect the sound with the knocks is just learning to recognize the right kind of sounds as knocks.<sup>16</sup>

It is, of course, true that generality is built into the abilities under consideration. The abilities involve being prone to respond to phenomena of a certain type in a certain way. It would be entirely natural for the deer-tracker to say that you can tell from the presence of tracks like these that deer have recently passed or to say that the presence of tracks like these is a sign (tells you) that deer have recently passed. In the light of this it might be thought that what I called the standard model of knowledge from indicators is fine at least so far as it goes. But, first, as already remarked, it is doubtful that in all cases in which it would be correct to ascribe the ability it would be correct to credit the subject with evidence adequate to establishing the truth of an appropriate generalization. Second, the generalizations that come to the mind of subjects and theorists alike are somewhat indeterminate because of the elusiveness of the relevant phenomena. Even in a simple case like that of telling that it has rained from the wetness of the streets and pavements we are hard put to characterize just what sort of wetness counts.

186

We know it when we see it, yet if we ask precisely what generalization we need to have support for in order to treat the wetness on some occasion as an indicator of recent rain, no very clear answer can be forthcoming. That, I think, is a sign that a sensibility is in play rather than a command, exercised at the level of judgement, of which features wetness must have to make it wetness of the rain-indicating sort. This is not to say that covering generalizations can play no role in our thinking. On the contrary, they form part of the understanding that informs our recognitional abilities and that we exploit when explaining and justifying the judgements we form. The crucial point though is that we account for the acquisition of knowledge in these cases in terms of the exercise of an ability to recognize a phenomenon as having a certain significance. It is the ability that is in the driving seat and its possession does not turn on independent support for any generalization that informs it.

Considerations analogous to those just set out apply to the epistemology of testimony. They bear on our ability to recognize speech acts as acts of telling. This ability is inculcated through our initiation into the practice of informing. But we should also take seriously the idea that our knowledge that  $p$  from someone's telling us that  $p$  is recognitional as well. It obviously involves recognition because we have to recognize the relevant utterance as an act of telling us that  $p$ . On the account I am proposing we should also think of our ability to tell that  $p$  from being told that  $p$  as being recognitional. The upshot of the exercise of such an ability is recognition that our interlocutor is telling us the truth. I exercise such an ability in telling from what my son says that my book is at home on my desk. Having such an ability I shall be prone to respond to utterances like this one in respect of content, manner of speaking, and so on, in circumstances like this, in the way I do on this occasion. While, stepping back, I can gesture towards features of the utterance that are truth-indicative, I need not register these at the level of judgement and it is accordingly implausible that my acceptance of a covering generalization capturing their significance figures as a premise in reasoning to the conclusion that the thing told is true.

What drives the account of knowing the truth in terms of recognizing the truth-indicative significance of the telling is not simply the phenomenological immediacy with which the truth is taken in. It is that while there are features of the act of telling that are undoubtedly relevant to our reception of the testimony, we do not register these at the level of judgement. The lesson to draw, I suggest, is that these responses have a greater kinship with perceptual-recognitional judgements than might at first seem. That is why we have to bring into play the notion of a sensibility not only to account for the identification of acts of telling but also to account for our taking in the truth-indicative significance of what we are told.<sup>17</sup>

#### 4. SENSIBILITY AND TRUSTWORTHINESS

What I have to say about the sensibility by which we recognize an act of telling to be truth-indicative merely gives a clue to the direction in which I think we should go in these matters. The issue merits further exploration. However, I doubt that philosophical reflection will yield anything like a theory as opposed to making it progressively more plausible that testimonial knowledge in the straightforward cases is best conceived in terms of the exercise of a sensibility and not in accordance with the standard way of thinking about knowledge from indicators.

We should note first that although the practice of informing by telling goes a considerable way towards explaining how it is that we so often acquire true beliefs from testimony it does not take us all of the way. The practice is in place only because often enough people follow its rules. Beyond mere participation in the practice there are additional incentives to avoid breaches of its rules: deception and incompetence are liable to be exposed, to the detriment of the informant. These factors explain why in a suitable setting even the gullible will often acquire true beliefs by accepting what they are told and the more



discriminating can do so with a fairly high degree of reliability simply by accepting what they are told when they do not detect any reason to do otherwise.<sup>18</sup> They impose limits on believing falsely from testimony but they do not illuminate what makes for knowledge or well-grounded belief in particular cases. For instance, they do not explain why someone's telling us something can settle the matter so that we acquire knowledge from what we are told. While it is clearly relevant to the explanation of our ability to gain knowledge in the straightforward cases that there is a practice of informing by telling, which conditions our responses to testimony, an account of how we gain knowledge must say something about why we readily trust particular individuals and why we can be right to do so. A natural thought at this point is that to acquire knowledge from a person in a straightforward case of telling, (a) the person has to be trustworthy, and (b) we

188

must have knowledge of the person that enables us to recognize him or her as trustworthy on the matter in hand on the occasion in question. Trustworthiness has two components, *sincerity* and *competence*.<sup>19</sup>

Competence with respect to whether  $p$  is knowing whether  $p$ , which entails knowing that  $p$  if  $p$ , and knowing that not- $p$  if not- $p$ . A crucial dimension of competence is being in command of the requisite ways of telling. My son meets this condition in my example because he knows the title of the book about which I am enquiring and can determine by looking whether a book of that title is in my study. Obviously, merely being in command of the requisite ways of telling is not enough. The ways of telling must have been exercised. My son meets this condition because he has seen that the book is on my study desk.

Sincerity, as I shall understand it, has three strands: you are sincere in telling me that  $p$  if and only if (i) you believe that  $p$ , (ii) you believe that you speak from knowledge, and (iii) you intend that I should come to know that  $p$  through your telling me. Condition (i) needs no explanation. The rationale of condition (ii) is provided by the nature of telling as an act by which you give it to be understood that you know. If you believe what you say but think that your standing on the matter is less than knowledge then you are misleading those you tell because your act of telling gives it to be understood that you know. The rationale for condition (iii) is similar. As I observed in section 2, when commenting on the (B)-rule for telling, there are cases when informants speak from knowledge but whose telling is infelicitous because they do not intend that the recipient should be informed by their telling. All that I am adding here is that this is a form of insincerity.

The practice of informing by telling imposes limits on the extent to which we are likely to believe falsely by accepting what we are told. Within the context provided by the practice, trustworthiness plays a role in relation to knowledge from being told that is analogous to the role played by favourable environments in relation to perceptual knowledge and knowledge from indicators. To tell of a structure that it is a barn from its visual appearance, the environment in which the barn is situated must be favourable to the exercise of the relevant recognitional ability: in that environment barns must have a

*distinctive appearance or set of such appearances.* An appearance is distinctive of barns if and only if not easily could something have this appearance and not be a barn. In a fake-barn scenario (as in Goldman 1976) barns do not have a distinctive appearance because too many of the structures that look like barns are not. No one could have the ability to tell, and thus come to know, of things in such an environment that they are barns from their visual appearance to the points of view that would normally suffice for this purpose. The conclusion to draw from this is that the ability in question is tied to—constitutively dependent for its exercise

189

on—favourable environments.<sup>20</sup> There are analogues of fake-barn scenarios for cases of telling from the occurrence of an indicator phenomenon. Think of a shepherd in an environment in which a certain bleating sound is distinctive of sheep: not easily could there be that sound in these parts and it not be produced by sheep.<sup>21</sup> With respect to such an environment it is possible to learn, as the shepherd has, to tell that there are sheep nearby from the sound of bleating. But suppose the shepherd to be transported to an environment superficially looking like his familiar environment but in which there are sheep that produce the familiar bleating sound and goats that produce an indistinguishable sound. In this environment the bleating sound is not distinctive of sheep. Because all too easily could something produce that sound and not be a sheep it would not be possible in that environment to tell that sheep are nearby just from the sound of bleating. Now, a trustworthy person with respect to a certain subject matter is one who is such that not easily could that person tell you something concerning that subject matter and speak falsely. Just as in suitable environments a suitably equipped person may legitimately take at face value the appearance of barns, or a bleating sound, so those who are suitably equipped may legitimately take at face value the acts of telling of a person who is trustworthy with respect to a certain subject matter.

If Bill is my friend and an expert gardener, I can know him well enough to be able to take what he tells me about gardening at face value. We shall not be able to do justice to the role of experience gained through acquaintance with a person if we stick with the model of accumulating evidence that we sift and weigh, and exploit in reasoning. Apart from anything else we do not keep a record of all of the evidence we glean, but nor can we pin down all the factors that are germane to our responses to what the person tells us. Yet it is not a mystery that our experience of a person can equip us with an ability to recognize the truth-indicative significance of an act of telling. We are not baffled by the idea that I come to know that the book about which I am enquiring is on my study desk at home from my son's telling me that it is. I come to know because I know him well enough to have acquired an ability to tell from certain of his acts of telling that what he is telling me in those acts of telling is true. In the scenario envisaged his competence on the matter of the book's location is not an issue. But nor is his sincerity. My past experience of him puts me in a position to respond to his report in an atmosphere of trust. But the role of past experience is that of attuning me to the significance of his acts of telling with respect to the truth of what he tells. This kind of role for experience is not a peculiarity of

testimonial cases. Experience also subtly shapes the abilities exercised in perceptual knowledge and in non-testimonial indicator cases.

190

## 5. OBJECTIONS BRIEFLY CONSIDERED

The suggested account is likely to provoke an objection on the grounds that often there can be knowledge from testimony when the recipient does not know the informant well enough to have the kind of recognitional ability illustrated in the case of my son's testimony. Yet cases in which this is so might be just like what I have been calling straightforward cases.

Consider, for instance, informants like ticket clerks, shop assistants, and representatives of organizations who inform us about their products. We usually do not know these people personally though we sometimes acquire knowledge from them in the way we do in the straightforward cases. It might be thought that there is nothing about them from which we can tell that they are trustworthy. Such a view overlooks the function of roles and practices. In addition to the practice of informing there are practices associated with information-disseminating roles. People who occupy such roles have little incentive to mislead and the risks to their doing so are significant. They will have acquired training so that they will generally be competent on the sort of matters on which we are likely to trust them. Providing that we are sufficiently discriminating about, for instance, the roles they occupy, and about the contents of what they tell us, we can be in a position to know some of the things they tell us. Here too experience hones our discriminative powers.

What about the stranger in the street whom one asks for directions to the castle?<sup>22</sup> Can we not obtain knowledge from such a person while not knowing him or her personally and knowing nothing of any information-disseminating role the person occupies? Perhaps. It depends on how the scenario is spelled out and it is wise not to prejudge cases that are under-described. But it would be no objection to my account if we do not acquire knowledge in such cases. What we gain from testimony will often be less than knowledge but useful nevertheless. Provided one exercises a modicum of discrimination it is entirely reasonable to act on the information we receive from a stranger when little is at stake or when one has no alternative. Reasonably acting on information supplied does not even require belief, far less knowledge. All it requires is that there should be a reasonable chance that what one is being told is true. That will be guaranteed by the fact that there is a practice of informing by telling, along with the fact that there is nothing suspicious about one's informant. This does not force us into the standard defaultist position, since it concerns the reasonableness of action rather than justified belief or knowledge.

A rather different line of objection focuses on early learning. When we are young we take on board a great deal of information from parents, teachers,

191

and others. We are in no position to have acquired the kind of recognitional abilities that are central to my account. If early learning is to be conceived as the acquisition of

knowledge through being told, as in the straightforward cases, then the knowledge will not meet the conditions I have laid down. But there is no reason to view early learning in this way. We acquire a picture of the world through early learning, but nothing compels us to treat the picture in its early stages as embodying knowledge-through-being-told as I have conceived it in relation to straightforward cases. Knowledge of the latter kind is knowledge that we gain on someone's say so. Its status as knowledge, all else equal, depends on the subject's continuing to have a reason to believe the thing told—a reason that is constituted by the fact that he or she has been told it. When we were young acts of telling undoubtedly contributed to inculcating facts in us but it does not follow that what we gained in this way is to be understood on the model of straightforward testimonial interchanges. Much of what we imbibe when young never becomes knowledge. Some of it, for instance, basic geographical and historical facts, becomes knowledge through repeated encounters with sources of the facts in question. But we often lose touch with the relevant sources in the sense that we could not say from which sources we acquired this or that item of knowledge. In these cases our knowing the fact in question is not like knowing that an ex-colleague was at some conference because a present colleague reports having met him there. In the latter case one knows not just because of, but *in view of*, having been told, and one has a reason to believe the thing told, which is supplied by the fact that one has been told it. However, it could be that I know, say, that Hobart is the state capital of Tasmania, without being able to pin down any particular source or sources from which I gleaned this information. I have a good idea of relevant kinds of source—teachers, atlases, encyclopedias, newspaper reports, and so forth. I also know how to confirm this item of information if called on to do so. Yet while I have been caused to believe that Hobart is the state capital of Tasmania by encounters with suitable sources, this is not a clear case of knowledge based on evidence, by which I mean knowledge the current standing of which as knowledge is explained by my currently being in possession of evidence that gives me a reason to believe that Hobart is the state capital of Tasmania. Rather, this knowledge consists in an ability to recall a publicly available, known fact, which has been gained from repeated encounters with reliable sources of information.<sup>23</sup> The role of past experience has been to expose me to reliable sources of information and to inculcate the ability to recall the facts that those sources made available. It

192

is not that of having supplied me with evidence in view of which I now take the fact in question to be a fact.<sup>24</sup> The upshot, I suggest, is that early learning should not be a test case for the adequacy of an account of straightforward cases of acquiring knowledge from testimony. I have no more than hinted at a line of approach to early learning that is compatible with the account I have proposed of the straightforward cases.<sup>25</sup>

193

## 10 Can A Priori Entitlement be Preserved by Testimony?

Ram Neta

Before we address the question that forms the title of this chapter—a question that Tyler Burge articulated and famously and controversially answered in the affirmative—let us begin by clarifying it.<sup>1</sup> Doing this will take considerable work, since we have to clarify what is meant by ‘entitlement’, what is meant by speaking of an entitlement’s being ‘a priori’, what is meant by speaking of this a priori entitlement being ‘preserved’, and finally, what is meant by speaking of its being so preserved ‘by testimony’. In the course of clarifying these terms, we will locate the resources necessary for Burge to defend his view from a prominent recent criticism, a criticism that focuses on the fact that Burge regards the entitlements preserved by testimony as typically a priori. But we will also find that Burge’s argument for his view can seem plausible to us only to the extent that we confuse entitlement with justification, and this is a confusion against which Burge himself warns.

I’ll begin by saying what Burge means by ‘entitlement’, then proceed to clarify what is meant by speaking of an entitlement’s being ‘a priori’. Next, I’ll discuss the notion of this a priori entitlement being ‘preserved’, and finally, I’ll say what it is for it to be preserved ‘by testimony’. Once all that is on the table, we’ll then turn to the task of defending Burge’s thesis from the criticism of it that we find in Malmgren (2006). This defence will help us to understand Burge’s view more clearly, and also help us to see that what’s implausible about Burge’s view has nothing in particular to do with his widely contested claims about the a prioricity of our testimonially preserved warrants.<sup>2</sup> (Indeed, Burge may very well be right in claiming that testimonially preserved warrants are sometimes a priori.) What is implausible about Burge’s view has to do rather with his generally uncontested claim that testimonially preserved warrants are sometimes *entitlements*.

194

### 1. ENTITLEMENT

So, first, what property does Burge mean to denote by the term ‘entitlement’? We can answer this question by sketching the ambitious theoretical picture within which the notion of entitlement is supposed to fit.

For Burge, animals, and their functional subsystems, have ends, goals, commitments, or functions. The satisfaction of one of these ends, the attainment of one of these goals, the fulfilment of one of these commitments, or the successful performance of one of these functions is a good for the creature, or the functional subsystem, of which it is an end, goal, commitment, or function. This is as true of the creature’s psychological subsystems as it is of the creature’s respiratory, circulatory, or reproductive subsystems: each such system has a function, and the successful performance of that function is a good for the functional subsystem. So, for example, a creature’s circulatory system functions to circulate nutrients to all of the creature’s cells, and the successful performance of that

function is a good for the circulatory system. (It is typically a good for the creature as well, but it need not be—for instance, the creature's overall health or survival may, in some cases, depend on the starvation of a particular part of the creature's body, and in that case what is a good for the circulatory system is not a good for the creature.) Another example: a creature's representational system functions to provide correct representations. And more specifically, a creature's belief system functions to provide true beliefs. While it may sometimes be good for a creature to have false beliefs, it is a good for the creature's belief system to have true beliefs.<sup>3</sup>

Now, given a creature's habitat and its perspectival and other representational limitations, there will be some ways of forming beliefs that are, necessarily and a priori, likely (though not certain) to lead to true beliefs. Beliefs that are so formed under normal circumstances in the creature's habitat will themselves be, necessarily and a priori, likely (though not certain) to be true. On Burge's view, what it is for a creature to be warranted in forming beliefs in a particular way, or for a particular belief to be *warranted*, is just for that way of forming beliefs, or for that particular belief, to fall into the respective categories just described—the categories that necessarily and a priori involve likelihood of truth. For a way of forming beliefs, or a particular belief, to enjoy the epistemic good warrant, is for it to be such that, in forming beliefs in that way, or in forming that particular belief, the creature is doing just what creatures so constituted *should* do, when under normal circumstances in their habitat, in order to attain true belief. The epistemic good of warrant is in this way related to the belief system's good of true belief.

195

This is how Burge understands warrant. But there are two kinds of warrant: entitlement and justification. Burge explains the difference between these two kinds of warrant in the following passage:

The distinction between justification and entitlement is this: Although both have positive force in rationally supporting a propositional attitude or cognitive practice, and in constituting an epistemic right to it, entitlements are epistemic rights or warrants that need not be understood by or even accessible to the subject. We are entitled to rely, other things equal, on perception, memory, deductive and inductive reasoning, and on—I will claim—the word of others. The unsophisticated are entitled to rely on their perceptual beliefs. Philosophers may articulate these entitlements. But being entitled does not require being able to justify reliance on these resources, or even to conceive such a justification. Justifications . . . involve reasons that people have and have access to. These may include self-sufficient premises or more discursive justifications. But they must be available in the cognitive repertoire of the subject. (Burge 1993: 458–9)

This passage strongly suggests that Burge intends the distinction between entitlement and justification to be understood as follows:

If C is a creature and E is a warranted epistemic act or state of C's, then C is justified in E if C can, upon reflection, *offer an articulate reason for E*. If C cannot do this, then C's warrant for E is an entitlement.

This interpretation of Burge's distinction is corroborated by other passages, such as this:

Entitlement is *epistemically externalist* inasmuch as it is warrant that need not be fully conceptually accessible, even on reflection, to the warranted individual. The individual need not have the concepts necessary to think the propositional content that formulates the warrant. . . . The other primary sub-species of epistemic warrant is *justification*. Justification is warrant by reason that is conceptually accessible on reflection to the warranted individual. (Burge 2003: 504)

Unfortunately for his reader, at other points, Burge's writing suggests—or at least may seem to suggest—a different distinction between entitlement and justification. For instance:

An individual's epistemic warrant may consist in a justification that the individual has for a belief or other epistemic act or state. But it may also be an *entitlement* that consists in a status of operating in an appropriate way in accord with norms of reason, even when these norms cannot be articulated by the individual who has that status. We have an entitlement to certain perceptual beliefs or to certain logical inferences even though we may lack reasons or justifications for them. The entitlement could in principle presumably—though often only with extreme philosophical difficulty—be articulated by someone. But this articulation need not be part of the repertoire of the individual that has the entitlement. (Burge 1996: 93–4)

In this passage, when contrasting entitlements with justifications, Burge says that entitlements consist 'in a status of operating in an appropriate way in accord with

196

norms of reason, even when these norms cannot be articulated by the individual who has that status'. This passage may seem to some readers to suggest that Burge's distinction between entitlements and justification should be understood as follows:

If C is a creature and E is a warranted epistemic act or state of C's, then C is justified in E if C can articulate the norms according to which E is warranted. If C cannot do this, then C's warrant for E is an entitlement.

One problem with this second interpretation of Burge's distinction between entitlement and justification is that, if this is how Burge intends to draw the distinction, then it makes it very hard to understand why Burge says what he does in the passages just quoted from Burge (1993) and Burge (2003). A second and perhaps more serious problem is that, if this is how to understand Burge's distinction between entitlement and justification, then justifications are possessed only by creatures that are capable of articulating the epistemic norms to which they are answerable. But, given his other commitments, Burge cannot reasonably restrict so narrowly the range of creatures that can enjoy justifications. For

instance, Burge tells us that only creatures that have propositional attitudes can enjoy entitlements, but only creatures that are capable of justifying their propositional attitudes are capable of having such attitudes.<sup>4</sup> Thus, if Burge were to claim that justifications could be enjoyed only by creatures that are capable of articulating the norms according to which their attitudes are warranted, he would be committed to claiming that entitlements could be enjoyed only by such creatures. Now, Burge clearly and repeatedly states that young children and some non-human animals can enjoy entitlements.<sup>5</sup> Therefore, if he were to claim that justifications could be enjoyed only by creatures that are capable of articulating the norms according to which their attitudes are warranted, he would be committed to the obviously false claim that young children and some non-human animals are capable of articulating the norms according to which their attitudes are warranted. We should therefore avoid attributing to Burge the view that justifications could be enjoyed only by creatures that are capable of articulating the norms according to which their attitudes are warranted.

Why, then, does Burge (1996) contrast entitlements with justifications by saying that entitlements consist 'in a status of operating in an appropriate way in accord with norms of reason, even when these norms cannot be articulated by

197

the individual who has that status'? I think this sentence should be understood not as suggesting that justification for believing that  $p$  requires the ability to articulate the norms of reason that make it appropriate for one to believe that  $p$ , but rather as suggesting that entitlement to believe that  $p$  requires *the lack* of such an ability. When a believer is able to articulate the norms by virtue of which she is warranted, then her possession of that ability is sufficient, *but not necessary*, to her warrant's being a justification: in such a case, the believer can offer an articulate reason for her warranted act or state, and she can do this by citing the norm by virtue of which it is warranted. Therefore, it is necessary, but not sufficient, for a warrant's being an entitlement that the believer cannot articulate the norms by virtue of which she enjoys that warrant. This means that it is sufficient, but not necessary, for a warrant's being a justification that the believer can articulate the norms by virtue of which she enjoys the warrant. On Burge's view, as I understand it, a warrant is a justification whenever the believer can articulate the warrant. In particular, if a believer is warranted in believing that  $p$  on the basis of an inference to her conclusion that  $p$ , then, for Burge, this is sufficient for the believer's warrant for  $p$  to be a justification.

We've now said all that we're going to say about what sort of property Burge means to denote by means of the term 'entitlement'. What makes an entitlement '*a priori*'?

## 2. A PRIORICITY

Whenever a creature  $C$  is warranted in an epistemic act or state  $E$ , there is something that *makes*  $C$  warranted in  $E$ . In other words, there is something that constitutes  $C$ 's warrant for  $E$ . Whatever this thing is that constitutes  $C$ 's warrant for  $E$ , either that thing includes



C's specific perceptual and sensory experiences, or it does not. If it does, then C's warrant for E is a posteriori (or, in other words, empirical). Otherwise, C's warrant for E is a priori. I should note that this conception of the a priori is not entirely standard today (not that there is any conception of the a priori that is entirely standard today!). It will be important to keep Burge's conception of a prioricity in mind throughout this paper, especially since failure to keep it in mind has occasioned some misunderstanding among his critics.

Where I've spoken of what constitutes someone's warrant, Burge speaks of what constitutes the 'justificational force' of a warrant (i.e. a justification or an entitlement)—in doing so, of course, Burge is using the term 'justificational force' more broadly than he uses the term 'justification' (which he intends to contrast with entitlement). Here's how Burge explains a prioricity: 'A justification or entitlement is a priori if its justificational force is in no way constituted or enhanced by reference to or reliance on the specifics of some range of sense

198

experiences or perceptual beliefs' (Burge 1993: 458). And he elaborates on this formulation:

An individual need not make reference to sense experiences for his justification or entitlement to be empirical. My term 'reliance on', in the explication of apriority, is meant to acknowledge that most perceptual beliefs about physical objects or properties do not refer to sense experiences or their perceptual content. Such beliefs make reference only to physical objects or properties. But the individual is empirically entitled to these perceptual beliefs. The justificational force of the entitlement backing such beliefs partly consists in the individual's having certain sense experiences, or at any rate in the individual's perceptual beliefs being perceptual. (Burge 1993: 460)

Notice that Burge's distinction between the a priori and the a posteriori is a distinction that applies to warrants in general, and it cuts across the distinction between entitlements and justification. So long as what makes C entitled to E does not include any of C's specific perceptual or sensory experiences, C's entitlement to E is a priori—although C might, of course, fail to know this, since she might not know, and not even have the ability to conceive of, what entitles her to E. (Recall that entitlements need not be understandable by those who enjoy them.) Therefore, C might have an a priori entitlement to E, even though C has no idea that she is so entitled, and is completely unable to offer any a priori justification for E. So, even if we are entirely unaware of the a priori entitlements that we enjoy by virtue of testimony, this fact does nothing to tell against our enjoyment of such a priori entitlements. Again, this point will become important later on, when we consider how Burge can defend his view against a recent criticism.

### 3. PRESERVING ENTITLEMENT

We've said what entitlements are, and what it is for an entitlement to be a priori. Now, what is it for an entitlement, or any other kind of warrant, to be *preserved*? Burge does not explicitly answer this question, but his remarks in Burge (1993)<sup>6</sup> suggest a partial answer. In criticizing Chisholm's conception of the role of memory in providing us with justification for believing the conclusions of long deductive arguments, Burge writes:

If memory supplied, as part of the demonstration, 'contingent propositions about what we happen to remember,' the demonstration could not be purely logical or mathematical. But the normal role of memory in demonstrative reasoning is, I think, different. Memory does not supply for the demonstration propositions about memory, the reasoner, or past events. It supplies the propositions that serve as links in the demonstration itself. Or

199

rather, it *preserves* them, together with their judgmental force, and makes them available for us at later times. Normally, the content of the knowledge of a longer demonstration is no more about memory, the reasoner, or contingent events than that of a shorter demonstration. One does not justify the demonstration by appeals to memory. One justifies it by appeals to the steps and the inferential transitions of the demonstration.

(Burge 1993: 462)

In this passage, Burge is contrasting two different ways in which memory can be involved in furnishing us with inferential justification: it can be involved by giving us justification for believing various contingent propositions about our memory, or it can be involved simply by preserving 'the propositions that serve as links in the demonstration itself' along with 'their judgmental force', and it 'makes them available for use at later times'. But what is it to preserve a proposition, or its judgemental force? Presumably, Burge means that our memory of the proposition can be preserved, along with our memory of the proposition's being (as we judge) true: just as I can now judge that  $68 + 57 = 125$ , I can later on remember that  $68 + 57 = 125$ . It's not simply that I will remember *the proposition* that  $68 + 57 = 125$ —rather, I will remember *that*  $68 + 57 = 125$ . What I remember—that  $68 + 57 = 125$ —has the same propositional content and the same judgemental force as my original judgement (i.e. the judgement that  $68 + 57 = 125$ ).

Besides preserving propositional content and judgemental force, memory also, according to Burge, makes propositions 'available for use at later times'. What does Burge mean when he speaks of such availability? The context of Burge's discussion suggests an answer: memory makes a proposition available for use at later times if it preserves one's warrant for the proposition, so that one continues to be in a position to draw warranted conclusions from the proposition. Right now, as I judge that  $68 + 57 = 125$ , I am justified in drawing the conclusion that  $125 - 68 = 57$ . But later on, as I remember that  $68 + 57 = 125$ , I will still be justified in drawing the conclusion that  $125 - 68 = 57$ : whatever warrant attached to my earlier judgement also attaches to the later judgement. The warrant is preserved, along with content and force, by memory.

This is one example Burge gives of warrant preservation. But, as we'll see in our discussion of Burge's views on the epistemology of testimony, it is not the only kind of warrant preservation that Burge countenances. Warrants (whether entitlements or justifications) can be preserved in at least two different ways: a warrant for a particular epistemic act or state can be preserved in memory from one moment to a later moment, and a person's warrant for a particular epistemic act or state can be preserved in testimony from one person to another. Inspection of Burge's various cases of warrant preservation reveals that warrant preservation involves three elements:

First, there are two distinct epistemic acts or states.

200

Second, one of these two epistemic acts or states occurs because the other one does.

Third, the former epistemic act or state is warranted because the latter is.

In sum, whenever warrant is preserved, one epistemic act or state occurs and is warranted because another epistemic act or state occurs and is warranted. I should make clear, by the way, that Burge does not tell us enough about warrant preservation to make clear whether the relation denoted by 'because' in the third of the three conditions above must be a *causal* relation: while warrant preservation involves a causal process, it's not clear that the relation between the warrants themselves is a causal relation.

I take the three conditions above to be necessary conditions of the phenomenon that Burge calls 'warrant preservation', but I do not know if they are jointly sufficient conditions. Burge does not say enough to allow us confidently to specify fully, and in a non-circular way, what is involved in warrant preservation. But for present purposes, this doesn't matter: we can make the points we need to make about Burge's view without attempting to enumerate conditions that are both necessary and sufficient for warrant preservation.

Just as entitlement is one form of warrant, so entitlement preservation is one form of warrant preservation. Entitlement preservation involves three elements:

First, there are two distinct epistemic acts or states.

Second, one of these two epistemic acts or states occurs because the other one does.

Third, the epistemic agent who is the locus of the former epistemic act or state is entitled to that act or state because the epistemic agent who is the locus of the latter epistemic act or state is warranted in that act or state. (Notice that it is not specified whether the former epistemic agent is identical to, or distinct from, the latter.)

And again, while these three elements are necessary for entitlement preservation, they may not be sufficient for it.

Notice that, in specifying the third condition of entitlement preservation, I describe the epistemic act or state that is at the receiving end of the preservative process as being one to which its agent enjoys entitlement, whereas I describe the epistemic act or state that is

at the initiating end of the preservative process as being one to which its agent enjoys warrant. The kind of testimonial process that we will be interested in discussing here is a process that begins with an epistemic act or state that is warranted somehow or other (perhaps the warrant takes the form of justification, perhaps it takes the form of entitlement), and ends with an epistemic act or state that is warranted in a way that its epistemic agent *need not be aware of*: in short, the sort of process that we will be interested in discussing here is a process that ends with an epistemic act or state to which its agent enjoys entitlement, and not (or not only) justification. Once again, this point will become important later on, when we consider how Burge can defend his view against a recent criticism.

#### 4. PRESERVING ENTITLEMENT BY TESTIMONY

I have said that entitlement preservation involves a causal process that begins with an epistemic act or state that is (somehow or other) warranted, and ends with an epistemic act or state to which its agent is warranted in a way that she need not be aware of, i.e. to which its agent enjoys entitlement. On Burge's view, there are at least two different kinds of causal process that sometimes have this profile. One of those processes is preservative memory, and the other—the one that will concern us here—is testimony. When Burge speaks of testimony, he means to denote an activity that may be, but is not necessarily, verbal—it is simply the activity of *presenting something as true*. If, while performing in a musical, I utter the declarative sentence ‘the rain in Spain falls mainly on the plain’, I have not thereby offered any testimony. If, while telling a joke, I begin by saying ‘A man walks into a bar’, once again I have not thereby offered any testimony. But if you ask me how many endowed chairs are in my department, and I say nothing but hold up five fingers, then, although I have uttered no words, I have (at least if my intentions and circumstances are otherwise normal) offered testimony: I have presented it as true that there are five endowed chairs in my department. If you ask me whether Jones has finally decided to pay his debts, then, even though you have asserted nothing in asking this question, you have nonetheless offered testimony: you have presented it as true that Jones has debts. While assertion is perhaps the most typical way to offer one's testimony, it is not the only way.<sup>7</sup>

Whenever a sincere act of testimony is understood and accepted by someone (call this latter agent the ‘auditor’, putting aside the misleading suggestion that testimony must be *heard* in order to be understood)<sup>8</sup>, there is a causal process of a certain kind running from the testifier to the auditor. This causal process starts with the epistemic act or state that the testifier was presenting as true in giving her testimony; and it ends with the epistemic act or state by virtue of which the auditor accepts the testimony. According to Burge, this causal process can, and often does, involve entitlement preservation. The epistemic act or state the content of which the testifier presents as true is warranted, and because it occurs and is warranted, the auditor is entitled to an epistemic act or state with the very same content and judgemental force.

Recall, however, that one may be warranted in believing what is completely false, and so one may be warranted in believing something on the basis of hallucinated testimony. Perhaps no rational source has ever presented it as true that, say, Ghana is in Africa. But it seems to me that a rational source has presented this to me as true, and I have no reason to doubt that a rational source has done so: in that case, I am warranted in believing that a rational source has presented it to me as true that Ghana is in Africa, and indeed, I am warranted even in believing that Ghana is in Africa, but no testimony to this effect has actually occurred. (In such a case, my belief, while true and warranted, would not constitute knowledge. This is a Gettier case.) My beliefs are warranted by my hallucination of testimony, but the warrant that I can enjoy by dint of this hallucination may be the same as the warrant that I can enjoy by dint of actual testimony, so long as there are no defeating reasons in either case. That, at any rate, is Burge's view.

Now that we've explained Burge's view, let's lay out his argument for this view.

## 5. BURGE'S ARGUMENT FOR THE VIEW THAT TESTIMONY PRESERVES A PRIORI ENTITLEMENT

Here is Burge's very compressed summary of his own argument:

We are a priori prima facie entitled to accept something that is prima facie intelligible and presented as true. For prima facie intelligible propositional contents prima facie presented as true bear an a priori prima facie conceptual relation to a rational source of true presentations-as-true: Intelligible propositional expressions presuppose rational abilities and entitlements; so intelligible presentations-as-true come prima facie backed by a rational source or resource for reason; and both the content of intelligible propositional presentations-as-true and the prima facie rationality of their source indicate a prima facie source of truth. (Burge 1993: 472)

This argument aims to establish the following conclusion a priori: if it seems to you that a particular proposition ( $p$ ) is presented as true by a rational source, then you are a priori prima facie<sup>9</sup> entitled to believe that  $p$ . (The passage that I've just quoted from Burge [1993] does not clearly make the qualification that I've just made concerning the entitlement that you receive by virtue of its *seeming to you* that  $p$  is presented as true by a rational source. I add the qualification that the entitlement is received by its so *seeming to you* in order to accommodate the point made on Burge's behalf at the end of the preceding section, i.e. that the warrant that one has on the basis of testimony is warrant that one would equally

well have even if one merely hallucinated that testimony. Thus, Burge writes, 'In the order of epistemic warrant, seeming understanding is the rational starting point' [1997: 31].)

What a priori argument does Burge take to establish this conclusion? The passage quoted just above is very compressed, but, given the context in which it is set, it is reasonable to suppose that the steps of Burge's argument can be laid out as follows. In stating this argument, I use the schematic variable ' $p$ ' in step 5 and the conclusion to range over all propositions our understanding of which does not require perceptual identification of some deictic element. These are the propositions our understanding of which is, in Burge's words, 'purely intellectual',<sup>10</sup> and it is only these propositions, according to Burge, a priori entitlement to which can be preserved through testimony. (That is because it is only these propositions, Burge thinks, which one can ever be a priori warranted in believing.) I hereby introduce the term 'purely conceptual propositions' to denote this class of propositions.

I should make clear, before stating this argument, that there is not a single step of this argument that is uncontroversial (or, indeed, anything less than highly controversial!). I will not attempt to defend the steps of this argument here, but only to set them out.

- (1) Necessarily, rational sources are, other things being equal, likely not to believe what is not true. (It is supposed to follow a priori from the nature of rationality and from the nature of belief that a rational believer tends, though perhaps with many exceptions, to believe only what is true. The kind of likelihood that Burge has in mind here, and in the steps of the argument below, is *objective*, not subjective.)
- (2) Necessarily, rational sources are, other things being equal, likely not to present as true something that they do not believe to be true. (It is supposed to follow a priori from the nature of rationality and from the nature of testimony—of the activity of presenting something as true—that a rational testifier tends, though perhaps with many exceptions, to present as true only what she believes to be true.)
- (3) Necessarily, the content of a presentation-as-true is constitutively determined by relations between content-bearing items and their referents in such a way that presentations as true are, other things being equal, likely to be true. (It is supposed to follow a priori from the nature of content that contents presented as true are, though perhaps with many exceptions, true.)

204

- (4) Necessarily, what is presented as true by a rational source is, other things being equal, likely to be true. (This is supposed to be supported a priori by 1, 2, and 3. Notice that, if the inference from 1, 2, and 3 to 4 is sound, then the inference from 1 and 2 to 4 should be sound as well: 3 looks to be redundant. But Burge thinks that 3 is also true, and provides a ground for 4 that is distinct from, and reinforces, the basis for 4 provided by 1 and 2.)
- (5) Necessarily, if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$  is presented as true by a rational source. (It is supposed to follow a priori from the nature of understanding that what we seem to understand to be presented as true by a rational source tends, though perhaps with

- (4) Necessarily, what is presented as true by a rational source is, other things being equal, likely to be true. (This is supposed to be supported a priori by 1, 2, and 3. Notice that, if the inference from 1, 2, and 3 to 4 is sound, then the inference from 1 and 2 to 4 should be sound as well: 3 looks to be redundant. But Burge thinks that 3 is also true, and provides a ground for 4 that is distinct from, and reinforces, the basis for 4 provided by 1 and 2.)

many exceptions, to have actually been presented as true by a rational source; or in other words, apparent understanding tends, though perhaps with many exceptions, to be genuine understanding. And if there is a necessarily reliable connection between apparent understanding of  $p$  as having been presented as true and  $p$ 's having actually been presented as true, then, given Burge's account of the nature of warrant, it follows that, if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$  is presented as true by a rational source.)

- (6) Necessarily, if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$ . (5 tells us that necessarily, if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$  is so presented. But if you are warranted in believing that  $p$  is presented as true, and if 4 is also true, then it follows that, necessarily, any process that leads a believer to whom it seems that  $p$  has been presented as true to believe that  $p$  will be a reliable process. And so, from the nature of warrant, it follows that, necessarily, if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$ .)

Burge takes each step of the argument above to be a priori *warranted*, and a priori knowable. But Burge's claims concerning what is a priori *warranted*, and a priori knowable, extend even farther than the claims in the argument above. For step 5 to be a priori is for it to be a priori that: If it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$  is presented as true by a rational source. But Burge wants to claim something over and above this. Burge wants to claim that, when it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that  $p$  is presented as true by a rational source. On what grounds can he make this stronger claim? Burge does not explicitly address this issue, but it seems to me that the most plausible account of his thinking on this issue takes it to go as follows:

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that it seems this way to you.

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing the conditional: if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$  is presented as true by a

rational source. (This second premise follows from the a prioricity of step 5, by disjunction introduction.)

---

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that you are warranted in believing that  $p$  is presented as true by a rational source.

If you are a priori warranted in believing that you are warranted in believing that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that  $p$  is presented as true by a rational source.

---

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that  $p$  is presented as true by a rational source. And if this warrant is not epistemically accessible to someone who enjoys it, then it is an entitlement.

Although this argument is obviously not formally valid, it still seems to have some plausibility, and it seems to me to provide the most plausible account of why Burge thinks that an auditor (or apparent auditor) has an a priori entitlement to believe that a rational source presented something as true. In any case, if this is not how Burge is thinking of the matter, then it is not at all clear on what grounds he could believe that an auditor (or apparent auditor) can have a priori entitlement to believe that a rational source presented something as true.

So that, I take it, is why Burge thinks that, if it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing (and typically, entitled to believe) that  $p$  is presented as true by a rational source. But now, why does Burge think that, if it seems to you that  $p$  (a purely conceptual proposition) is presented as true by a rational source, then you are a priori warranted in believing that  $p$ ? For step 6 of the argument above to be a priori is for it to be a priori that: If it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$ . But Burge wants to claim something over and above this. He wants to claim that, when it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that  $p$ . On what grounds can he claim this? Burge does not explicitly address this issue, but, again, it seems to me that the most plausible account of his thinking on this issue takes it to go as follows:

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that it seems this way to you.

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing the conditional: if it seems to you that  $p$  is presented as true by a rational source, then you are warranted in believing that  $p$ . (This second premise



follows from the a prioricity of the conclusion of the six-step argument set out above, by disjunction introduction.)

---

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that you are warranted in believing that  $p$ .

If you are a priori warranted in believing that you are warranted in believing that  $p$ , then you are a priori warranted in believing that  $p$ .

---

If it seems to you that  $p$  is presented as true by a rational source, then you are a priori warranted in believing that  $p$ . If this warrant is not epistemically accessible to someone who enjoys it, then it is an entitlement.

Again, although this inference is not formally valid, it still seems to have some plausibility, and this inference seems to me to provide the most plausible account of why Burge thinks that an auditor (or apparent auditor) has an a priori entitlement to believe the proposition that a rational source seems to have presented as true. In any case, if this is not how Burge is thinking of the matter, then it is not at all clear on what grounds he could believe that an auditor (or apparent auditor) can have this a priori entitlement.

Notice that, in attempting to give a plausible account of Burge's argument, I have interpreted Burge as accepting two instances of the following form of argument:

If  $p$ , then S is a priori warranted in believing that  $p$

If  $p$ , then S is a priori warranted in believing that: if  $p$  then S is warranted in believing that  $q$ .

---

If  $p$ , then S is a priori warranted in believing that S is warranted in believing that  $q$ .

If S is a priori warranted in believing that S is warranted in believing that  $q$ , then S is a priori warranted in believing that  $q$ .

---

If  $p$ , then S is a priori warranted in believing that  $q$ .

This form of argument, while not formally valid, does appear quite plausible. (To avert a possible misunderstanding: notice that, in claiming that, if S is a priori warranted in believing that S is warranted in believing that  $q$ , then S is a priori warranted in believing that  $q$ , I am not thereby committed to a controversial claim of the form that, if S is a priori warranted in believing that S is warranted in believing that  $q$ , then that very warrant [the one that S is a priori warranted in believing that she has] is a priori. Perhaps the warrant that S is a priori warranted in believing herself to have is not itself a priori, but if she is a priori warranted in believing that she has some such warrant, then she also has an a priori warrant.) We will eventually come back to the issue of why arguments

of this form appear plausible: in the concluding section of this chapter, I will argue that such arguments are compelling if the kind of warrant being discussed is justification, but not if it is entitlement.

Of course, Burge recognizes that there is not a single premise or inference in any of the arguments above that is uncontroversial, but his efforts to defend various premises and inferences of these arguments are scattered over numerous writings, and we cannot attempt to survey these efforts now. So I will not attempt here to assess any of Burge's controversial arguments for those premises. (I should note that I have not found, anywhere in Burge's *oeuvre*, a defence of his inference from 1, 2, and 3 to 4. I am not sure why he accepts this inference, but I also have no grounds on which to question it.) Rather, I will defend Burge's view against one prominent recent criticism of it. And finally, I will show that, even if we concede all of the substantive, controversial metaphysical claims that Burge makes in his argument 1—6 above, there is nonetheless another problem with Burge's argument, a problem that undercuts his argument for the claim that a priori entitlement can be preserved by testimony.

## 6. MALMGREN'S CRITICISM

Malmgren (2006) argues that testimony cannot possibly preserve a priori entitlement (or indeed, a priori warrant of any kind). Malmgren takes the problem with Burge's view to be his claim that the entitlement preserved by testimony is a priori: all testimonially preserved warrants must be a posteriori, according to Malmgren (and her view concerning this issue is one version of the Humean orthodoxy concerning testimony that Burge set out to challenge).

In this section, I will respond, on Burge's behalf, to Malmgren's arguments. More specifically, I will first respond to her argument that, for any purely conceptual proposition  $p$ , one cannot be a priori entitled to believe that a rational source presented it as true that  $p$ . Then I will respond to her argument that, for any purely conceptual proposition  $p$ , one cannot be a priori entitled to believe that  $p$  by virtue of receiving testimony that  $p$ .

So first, why does Malmgren think that, for any purely conceptual proposition  $p$ , one cannot be a priori entitled to believe that a rational source presented it as true that  $p$ ? Her argument is contained in this passage, quoted at length from her paper:

Burge . . . does seem to hold that the relevant entitlement is not just an entitlement to make a transition, but an outright entitlement to believe something—although something less specific than, say, that NN said that  $p$  at  $t$ . We are a priori entitled to believe that 'a given content is presented as true' (say, is asserted), where this is supposed to carry no

commitment to the content's being thus presented at any particular time or place, or by any particular person. We are a priori entitled to believe that it is said that  $p$ —period.

. . . Let us grant that we, normal recipients of testimony, do in fact have ‘non-committal’ beliefs of the relevant sort. . . . Suppose, further, that we have warrant for them. Is this warrant a priori warrant? It does not seem plausible to suppose that it is. Presumably, those beliefs typically rest—in both the psychological and epistemic sense of ‘rest’—on beliefs about what is said of the more specific type. But if that is correct, then the less specific beliefs are not a priori warranted, since the more specific beliefs are not a priori warranted.

Indeed, how could the less specific beliefs be a priori warranted? The proposition that it is said that  $p$  is highly contingent (just like the proposition that NN said that  $p$  at  $t$  is highly contingent). Arguably, there are a few examples of the contingent a priori. Perhaps I can know a priori that I am here now; and perhaps, if I stipulate that ‘Julius’ is to refer to the inventor of the zipper, I can know a priori that, if anyone uniquely invented the zipper, then Julius invented the zipper. But these examples are at best rare exceptions, exceptions that, on the face of it, have little in common with the case at hand. And the explanation(s) of how the contingent propositions in these examples can be known a priori do not seem to carry over. We cannot explain my allegedly a priori warrant for believing that it is said that  $p$  with appeal to the peculiarities of indexicals or of reference-fixing stipulations. How then do we explain it? (Malmgren 2006: 215–16)

Now, what I’d like to focus on is Malmgren’s suggestion in the final paragraph just quoted that we have no way to explain how it is possible for someone to whom it seems that a rational source has presented it as true that  $p$  to be a priori warranted in believing the contingent proposition that a rational source has presented it as true that  $p$ . It seems to me that, on the contrary, we have a perfectly good explanation of this, which I’ll now outline by means of an example.

Suppose that I present it as true that, say, Washington, DC is the capital of the United States. If I do this under normal circumstances, then, while I might not be a priori warranted in any particular beliefs concerning the medium by which I presented this proposition as true (did I assert it out loud, or did I express it by means of semaphores?), I can at least be a priori, defeasibly warranted in thinking that I have (somehow or other) presented it as true that Washington, DC is the capital of the United States: my warrant for believing this contingent proposition concerning what I have presented as true is of the same sort as my warrant for believing other contingent propositions concerning my communicative intentions, and it is not typically constituted by my sensory or perceptual experiences. It is crucial to recall here that for an agent to present something as true does not require that the agent make any noises or bodily movements: Burge leaves it open that an agent can present something as true by means of extra-sensory perception. Indeed, in at least one passage he suggests that judging that  $p$  is a way of presenting it as true that  $p$ : ‘I use the term “presentation as true” to cover more than assertions and *judgments*’ (Burge 1993: 482, emphasis)

added). If this is right, then I can present something to myself as true simply by thinking of it *as true*, that is, judging that it is true or otherwise affirming its truth in thought.

Now, if I can be a priori warranted in thinking that I have presented this proposition as true, then I can also be a priori warranted in thinking that I have *been given to understand* this proposition as true. But the following conditional is a priori: if I have been given to understand something as true, then a rational source (maybe me, maybe another rational source) has given it to be understood as true, or in other words, has presented it as true. And now, using the same form of reasoning that I attributed to Burge above, we can plausibly (albeit, not validly) reason as follows:

If it seems to me that I have been given to understand something as true, then I am a priori warranted in thinking that I have been given to understand something as true.

If it seems to me that I have been given to understand something as true, then I am a priori warranted in believing that, if it seems to me that I have been given to understand something as true, then a rational source has presented it as true.

---

If it seems to me that I have been given to understand something as true, then I am a priori warranted in believing that a rational source has presented it as true.

Although this inference is not valid, it is nonetheless plausible, and it seems to me to provide precisely the kind of explanation that Malmgren suggests cannot be provided. Thus, Malmgren's argument in the passage just quoted is unsuccessful. This is not for a moment to deny that, when it seems to me that I have been given to understand something as true, then I am a posteriori warranted in believing that a rational source has presented it as true: Burge's view leaves it wide open that, under such circumstances, we always have such a posteriori warrants. All that Burge is concerned to argue is that, under such circumstances, we *also* have the relevant a priori warrants.

Let's now turn to Malmgren's argument for the claim that, for any purely conceptual proposition  $p$ , one cannot be a priori entitled to believe that  $p$  by virtue of receiving testimony that  $p$ . Here's what Malmgren says:

Suppose John tells you that it is raining. Suppose furthermore that you have no reason to distrust him, no reason to distrust your senses, no reason to believe that you failed to understand his utterance (and so on). On the current view, are we not committed to saying that you have a priori warrant for the belief that it is raining? Surely this is absurd. (Malmgren 2006: 220)

One problem with this putative counterexample to Burge's thesis is that the proposition 'it is raining' is clearly not a purely conceptual proposition: our understanding of that proposition is not purely intellectual, but involves perceptual identification of a deictic reference to a particular place at which it is, according to the proposition, raining. So it is not a consequence of Burge's view that anyone can, under any circumstances, have a priori warrant for the belief that it is raining.

In response to this objection, Malmgren might alter her counterexample so that it involves a purely conceptual proposition. Burge offers an example of one such proposition: 'zebras are larger than red poppies'.<sup>11</sup> Isn't it implausible, Malmgren might ask, to think that anyone could, by virtue of having received some testimony, have a priori warrant for believing that zebras are larger than red poppies?

Well, what warrants you in believing that zebras are larger than red poppies? You *might* be warranted in believing it by having seen lots of zebras, and having seen lots of red poppies, and compared their visible sizes: in that case, your warrant would certainly not be a priori. But, whether or not you are warranted in this way, you might also be warranted in a different way: you might simply have been informed (somehow or other by means of testimony, the receipt of which you can no longer recall) of the sizes of zebras, and of the sizes of red poppies. Perhaps your receipt of this information required you to have certain perceptual experiences; but it is plausibly the information itself—and not the perceptual experiences by means of which you happened to receive it—that warrants your belief that zebras are larger than red poppies. (You could have received information by visual means, auditory means, or perhaps by means of ESP: it's not the medium, but rather your receipt of the message, that warrants your belief.) If you do have this kind of warrant for believing that zebras are larger than red poppies then, on Burge's broad conception of the a priori, you have an a priori warrant for believing that zebras are larger than red poppies. It's not obvious what, if anything, is implausible about this.

Malmgren has given us no good reason to doubt Burge's argument. So finally, let us turn to the task of evaluating Burge's argument ourselves.

## 7. WHAT'S WRONG WITH BURGE'S ARGUMENT HAS NOTHING TO DO WITH A PRIORICITY

Recall that, in reconstructing Burge's argument above, we interpreted Burge as making two instances of the following form of argument:

If  $p$ , then  $S$  is a priori warranted in believing that  $p$ .

If  $p$ , then  $S$  is a priori warranted in believing that: if  $p$  then  $S$  is warranted in believing that  $q$ .

---

If  $p$ , then  $S$  is a priori warranted in believing that  $S$  is warranted in believing that  $q$ .

If  $S$  is a priori warranted in believing that  $S$  is warranted in believing that  $q$ , then  $S$  is a priori warranted in believing that  $q$ .

---

If  $p$ , then  $S$  is a priori warranted in believing that  $q$ .

It is not easy to see what argument Burge could be giving to support his controversial claims concerning the a priori entitlements preserved by testimony unless he relies on the two arguments of this form that I set out in section 5. I take it then, that the success of Burge's arguments requires that this form of argument—or at least the two instances that Burge employs—be truth-preserving. Are they?

They may seem to be truth-preserving. To deny this would require either denying that a priori warrant is closed under modus ponens consequence, or else denying that having a priori warrant for believing that you have warrant for believing that  $p$  does not imply having a priori warrant for believing that  $p$ —and it may seem implausible to deny either of these things. Indeed, it's the plausibility of all this that makes it reasonable to interpret Burge as tacitly accepting arguments of this form for his controversial views concerning the a priori entitlements that are preserved by testimony.

But notice that Burge is claiming not simply that a priori warrants are preserved by testimony, but more specifically that a priori *entitlements* are preserved by testimony. So, what Burge needs is not just that some arguments of the form set out above be truth-preserving, but that those truth-preserving arguments justify a conclusion that is specifically about entitlement. Now, since the kind of a priori warrant that is preserved by testimony is typically, Burge thinks, one to which the auditor does not have epistemic access, can't we use arguments of the form set out above to show that the auditor has certain a priori warrants, and then conjoin these conclusions with the premise that the auditor has no epistemic access to these warrants to reach the conclusion that the auditor has a priori entitlements?

No. For the conditions under which arguments of the form set out above are truth-preserving might also be conditions under which the warrants mentioned in those arguments are warrants to which the auditor has epistemic access. In other words, we have not considered the possibility that arguments of the form set out above are truth-preserving only when the warrants mentioned in those arguments are justifications, and not entitlements.

Let's see whether this might be the case. It is plausible that what you're justified in believing is closed under modus ponens consequence, and it is also plausible that having an a priori justification to believe that you're justified in believing  $p$  suffices for having an a priori justification to believe that  $p$ . So, pending a defence of the controversial steps of his 1—6 argument given in section 5, Burge can, I think, plausibly argue that a priori justification is preserved by testimony. But what he claims is that a priori entitlement is preserved by testimony. To argue for this along the lines that I've offered him here, he'd have to give us some reason to believe two things:

212

- (a) what you're entitled to believe is closed under modus ponens consequence, and
- (b) having an a priori entitlement to believe that you're entitled to believe that  $p$  suffices for having an a priori entitlement to believe that  $p$ .

But, in fact, when we recall what Burge has to tell us about the nature of entitlement, it becomes clear that we have no reason to accept either (a) or (b).

Recall from our discussion of warrant in section 1 above that Burge thinks of warrants generally as attaching to beliefs by virtue of the process that leads to those beliefs—specifically, by virtue of the necessary reliability of that process under normal circumstances in the creature's habitat, given the creature's normal perspectival limitations. To say that a believer is *warranted* in believing that  $p$  is to say that there is, in the believer's present circumstances, some suitably reliable process that does or could lead the creature to believe that  $p$ . To say that a believer is *justified* in believing that  $p$  is to say that there is such a process, and that the creature has epistemic access to it as what warrants him in believing that  $p$ . And finally, to say that a believer is *entitled* to believe that  $p$  is to say that there is such a process, and that the creature does not have epistemic access to it as what warrants him in believing that  $p$ .

Now, justification is plausibly closed under modus ponens consequence for a simple reason. When there is a suitably reliable process that does or could lead a creature to believe that  $p$ , and a suitably reliable process that does or could lead a creature to believe that if  $p$ , then  $q$ , then there is bound to be a suitably reliable process that does or could lead a creature to believe that  $q$ : namely, deduction. But when a creature deduces  $q$  from the conjunction of  $p$  and (if  $p$ , then  $q$ ), then the creature's warrant for believing that  $q$  consists in this inference, and such inferential warrants, recall, are all justifications, for Burge. Thus, when a creature has a warrant for believing that  $p$ , and also has a warrant for believing that (if  $p$ , then  $q$ ), then that creature is bound to have a *justification* for believing that  $q$  (whether or not she avails herself of that justification and proceeds to believe that  $q$ ).

But why think that, under such conditions, the creature would also have an entitlement to believe that  $q$ ? For the creature to have such an entitlement, there would have to be another process, *distinct from deduction*, and one to which the creature does not have epistemic access, that does or could, under normal circumstances, lead the creature to believe that  $q$  whenever she believes both  $p$  and (if  $p$ , then  $q$ ). Why should we suppose that there is any such process? Burge certainly offers us no such reason, and I am aware of no such reason. So, even if justification is preserved under modus ponens consequence, we have no reason to suppose that the same is true of entitlement. In other words, we have no reason to accept (a) above.

Now consider (b). If a creature is justified in believing that she is justified in believing that  $p$ , then there is a suitably reliable process that does or could lead the creature to believe that  $p$ : namely, a suitably reliable inference from the premise that she's justified in believing that  $p$  to the conclusion that  $p$ . (If the creature is warranted in believing that  $p$ , then she is in a position to make a reliable, though of course not deductive, inference to  $p$ .) But notice, again, that if a creature infers  $p$  from the fact that she is warranted in believing that  $p$ , then the creature's warrant for believing that  $p$  consists in this inference,

and such inferential warrants are, once again, justifications. And so it follows that the creature has a justification for believing that  $p$ . So, if a creature is justified in believing that she's justified in believing that  $p$ , then it follows that she's justified in believing that  $p$ . But is the same sort of principle true of entitlement? Once again, only if there is some other process, *distinct from inference*, and one to which the creature does not have epistemic access, that does or could, under normal circumstances, lead the creature to believe that  $p$  whenever she believes that she's warranted in believing that  $p$ . Why should we suppose that there is any such process? Burge certainly offers us no such reason, and I am aware of no such reason. So, even if it's plausible that being justified in believing that you're justified in believing that  $p$  implies that you're justified in believing that  $p$ , there is still no reason to think that being entitled to believe that you're entitled to believe that  $p$  implies that you're entitled to believe that  $p$ . In short, we have no reason to accept (b) above.

I conclude that—whatever the merits of Burge's metaphysical views connecting rationality, intelligibility, truth, and presentation-as-true—Burge still gives us no reason to accept either (a) or (b) above, and so gives us no reason to accept crucial steps in his argument for the claim that a priori entitlements are preserved by testimony. Even if we grant Burge all of the controversial metaphysical claims in his argument 1—6 above, the farthest that this will get him is the claim that, if a rational source presents it as true that  $p$ , then you are entitled to believe *that a rational source has presented it as true that  $p$* , and also, if a rational source presents it as true that  $p$ , then you are entitled to believe that, *if a rational source has presented it as true that  $p$ , then you are warranted in believing that  $p$* . But we have not yet seen how these claims can jointly entail anything about what you're entitled to believe about the truth value of  $p$ , when a rational source presents it as true that  $p$ . It may seem that they do jointly entail that, if a rational source has presented it as true that  $p$ , then you are entitled to believe that  $p$ —but if it seems this way to us, then that is probably because we are confusing entitlement with justification.

214

The problem with Burge's argument is not that it fails to show that a priori entitlements are preserved by testimony, but rather that it fails to show that entitlements of any kind (as Burge conceives of entitlements) are preserved by testimony.<sup>12</sup>

215



## 11 The Assurance View of Testimony

Frederick F. Schmitt

### VERSIONS OF THE ASSURANCE VIEW

It is uncontroversial that testimony that  $p$  can give the hearer of the testimony an epistemic reason to believe  $p$ , and indeed make the hearer epistemically justified in believing  $p$ .<sup>1</sup> Call *evidentialism* the view that testimony gives such a reason by giving evidence that  $p$ , or at any rate by allowing the hearer justifiedly to take it that there is evidence that  $p$ . In its simplest version, evidentialism is a *reductive individualist* view of testimonial justification: the testimony gives such a reason because the hearer can infer that  $p$  from his or her belief that the testimonial act indicates that  $p$ , and the latter belief is justified either inductively or by a priori reasoning from a principle of charity in interpretation.<sup>2</sup> This view is *reductive* in understanding the reasons given in testimony to be based on the hearer's non-testimonially (in this case, inductively or a priori) justified beliefs. And it is *individualist* in understanding the giving of testimonial reasons to depend only on the belief  $p$ 's being based on epistemically justified beliefs of the hearer.

Contrasting with such a reductive view are *non-reductive* and *non-individualist* views.<sup>3</sup> I focus in this chapter on one non-reductive and non-individualist

216

view—the *assurance* view of testimony.<sup>4</sup> On this view, testimonial reason-giving or justification depends in some instances on the testifier's speech act of assuring the addressee that  $p$ , or on the addressee's acceptance of that assurance. Assurance allows an epistemic reason to be given to the addressee even when no evidence is given. The testifier's assurance may be understood in various ways—as the testifier's presenting himself or herself as having an adequate epistemic reason to believe  $p$  or as epistemically responsible in asserting  $p$ , or as inviting the addressee to trust him or her that  $p$ , or as offering to take responsibility for  $p$ 's being true. Given one or another understanding of assurance, the assurance view then says:

In some instances of testimony, the testifier assures that  $p$  and the addressee treats the testifier's speech act as assurance (rather than evidence), and the testimony gives the addressee an epistemic reason to believe  $p$  either because the testifier assures that  $p$  or because the addressee accepts that assurance (or because both of these conditions obtain).

It would be natural for a proponent of the assurance view to say all of the following:

the testifier's assurance or the addressee's acceptance of it makes available to the addressee a reason to believe  $p$ ;

the testifier's assurance or the addressee's acceptance of it gives the addressee an epistemic reason to believe  $p$ ;

the testifier's assurance or the addressee's acceptance of it makes available to the addressee a reason to believe *p*;

this giving of an epistemic reason entails that the addressee has an epistemic reason to believe *p*;

under certain conditions, the addressee's basing his or her belief *p* on this epistemic reason suffices for the belief's being testimonially justified.

I take it that the availability of the reason is not enough for the reason to be given, but the reason's being given *is* enough for the addressee to *have* the reason. Versions of the assurance view will differ on whether the addressee must *accept* the assurance in order for the testimony to give and the addressee to have the epistemic reason.<sup>5</sup> For the addressee to be testimonially justified in believing in the 'propositional' sense of justification, the epistemic reason need only be given and the addressee need only *have* the reason (assuming the reason given is good

217

enough for justification); versions of the assurance view will differ on whether that requires the addressee's acceptance of the assurance. For the addressee's belief to be testimonially justified in the 'doxastic' sense, the addressee must *base* that belief on the epistemic reason; this presumably requires that the belief arises from the addressee's acceptance of the assurance.<sup>6</sup>

Presumably merely giving an epistemic reason is not yet enough for giving a *good* epistemic reason or one *adequate* for the addressee's belief to be justified, or even for giving a reason having any positive quality.<sup>7</sup> Accordingly, Richard Moran adopts a two-tiered assurance view in his rich defense of the assurance view, 'Getting Told and Being Believed'.<sup>8</sup> In the first tier, assurance or its acceptance is enough for giving the recipient an epistemic reason to believe *p*, though not enough for giving the recipient a good epistemic reason to believe *p*.<sup>9</sup> In the second tier, the presence or absence of background conditions determines whether the reason given is good or bad:<sup>10</sup>

Whether this [the assurance or its acceptance] counts as a good or sufficient reason is not a matter of the speaker's illocutionary authority, but will depend both on his sincerity

218

and on his having discharged his epistemic responsibilities with respect to the belief in question. (295)

Assurance, although not sufficient for giving a good or bad reason, is necessary for the given reason even to count as epistemic and thereby licence the background conditions as determinants of whether the reason is good or bad:

These background conditions can themselves be construed as evidential, or at any rate not at the behest of the speaker to determine, but they are not themselves sufficient for giving any epistemic significance to the speaker's words, for the relevance of these conditions only comes into play once it is understood that a particular speech act is being performed with those words . . . The speaker has to constitute his utterance as having this or that illocutionary force before the empirical background conditions can contribute anything to its epistemic significance . . . Determining his utterance as an assertion is what gets the speaker's words into the realm of epistemic assessment in the first place. (289)

Moran here clarifies the assurance view by contrasting it explicitly with evidentialism and implicitly with alternative non-evidentialist views like the reliability, social responsibility, and transindividual reasons views.<sup>11</sup> On these alternative views, conditions like the testifier's sincerity, responsibility, knowledge, trustworthiness, and reliability in asserting  $p$  are sufficient for the utterance to give a good epistemic reason, and their absence is sufficient for the utterance to give a bad epistemic reason, so long as the right vehicle of testimony is in place. And the right vehicle may be something other than assurance—mere assertion, for example. On these views, what vehicle is right for testimony to give any epistemic reason (good or bad) is determined by what vehicle is appropriate for giving a good or bad epistemic reason as conceived by the view. Moran opposes these views by insisting that assurance is necessary for giving a good or bad epistemic reason in non-evidential cases, and this fact is independent of any good- or bad-making characteristics of the reason given by assurance. I take the two-tiered assurance view to be this:

- (i) Assurance that  $p$  or its acceptance is enough to give the recipient some epistemic reason to believe  $p$ .<sup>12</sup>
  - (ii) This reason is determined to be either good or bad by the presence or absence of 'background conditions'.<sup>13</sup>
- 219
- (iii) If the reasons are determined to be good, then the subject is testimonially justified in believing  $p$ .

The background conditions include some or all of the following: the testifier's sincerity in asserting  $p$ , 'his having discharged his epistemic responsibilities with respect to the belief in question' (Moran 2006: 295), the testifier's knowledge that  $p$ , trustworthiness, or reliability as to whether  $p$ , or at least the addressee's assumption that some or all of these conditions obtain (Moran 2006: 289).<sup>14</sup> I allow that all this is only one option for interpreting Moran, but I know of no better option.<sup>15</sup>

The assurance view divides into what I will call an *agreement* version and a *communicative norm* version. On the agreement version, assurance or its acceptance gives the addressee an epistemic reason to believe  $p$  in something like the way a promise and its acceptance give the promisee a practical reason to rely on the promised action. When assurance is accepted, this establishes an agreement of sorts between the testifier

and the addressee. The testifier initiates the agreement by presenting an epistemic reason to believe, inviting trust, or offering to assume responsibility for the truth of the proposition (conditional on the addressee's acceptance of that offer). The addressee may complete the agreement by accepting these. A completed agreement of this sort forms a personal relation between the testifier and the addressee. Reference to such a personal relation in turn explains the sense in which in testimonial uptake the addressee may be said to believe the testifier that  $p$ , and not merely believe the testifier's utterance.

The communicative norm version of the assurance view differs from the agreement version in treating assurance as giving the addressee an epistemic reason to believe in virtue of the fact that assurance is governed by norms of communication. These norms make demands on the activity of assuring and confer prima facie permission for any addressee to trust what is assured whether

220

or not an agreement is proffered or accepted.<sup>16</sup> This version of the assurance view does not entail that testimony gives an epistemic reason to believe in virtue of any actual personal relation between the testifier and the addressee or any reason to form such a relation. I believe that the communicative norm version faces fewer difficulties than the agreement version, but as no one has endorsed it, and Moran's view is clearly an agreement assurance view, I will focus on the agreement version and set aside the communicative norm version (only pausing later to note that the latter faces at least one difficulty that afflicts the agreement version). For convenience, I will refer to the agreement version simply as 'the assurance view'.

Turning now to two main variants of the agreement version, on one, a testimonial epistemic reason is given only by an *actual* agreement. Such an agreement generates commitments of both the testifier and the addressee, and these commitments give the recipient an epistemic reason to believe. Call this the *actual agreement* variant. On the other variant, the assurance by itself gives the addressee an epistemic reason to believe without any acceptance, hence without an agreement. It does so because in virtue of its content it *offers the recipient a benefit* great enough to give a practical reason to accept the assurance and fulfil its terms. Call this the *offer-of-benefit* variant. The difference between these two variants is analogous to the difference between the reason one has to fulfil the terms of a contract simply because one has made the contract and thereby agreed to fulfil its terms, and the reason one has to sign a contract and fulfil its terms because of the benefit contingent on signing. Moran's insistence that reason-giving depends on a personal relation might suggest attributing to him the actual agreement variant. But his core account of reason-giving is in fact an offer-of-benefit variant, and I believe this is the better choice for defending the assurance view.<sup>17</sup>

221

I grant at the outset that there is an important and pervasive phenomenon of assurance, understood to involve some or all of the following: presenting oneself as having adequate

epistemic reason to believe *p*, presenting oneself as being epistemically responsible in asserting *p*, inviting the addressee to trust that *p*, or offering to assume responsibility for *p*'s being true. And when assurance involves the latter speech acts, the addressee often enough accepts it. I grant too that assurance or its acceptance can give the addressee some *practical* reason to trust the testifier beyond what a hearer has in a case of mere assertion. My question is rather whether assurance or its acceptance gives the addressee an *epistemic* reason to believe, in the sense of a reason that contributes to the addressee's epistemic justification for believing *p*.

I will begin by criticizing the putative analogy between assuring and promising. I will then object to the assurance view, on three grounds. First, though assurance may give the recipient a practical reason to trust the testifier as to whether *p*, neither assurance nor its acceptance is *sufficient* to give the addressee an epistemic reason to believe *p*. Second, assurance isn't *necessary* for testimony to give the addressee a non-evidential epistemic reason to believe the proposition. Third, the assurance view cannot account for the conferring of epistemic reasons by testimony in the most important class of cases for a non-reductive view, the *primitive* ones for which the hearer lacks a substantial base of non-testimonial reasons. My first and second objections are consonant with those made by Jennifer Lackey (2008: ch. 8), to which I am much indebted here, though my arguments generally differ from hers. At the end of the paper, I will suggest an alternative non-reductive, non-individualist view of testimonial justification based on speech acts.

#### PRELIMINARY: ASSURANCE AND ASSERTION

The assurance view, as Moran presents it, assumes that assurance involves something beyond mere assertion in virtue of which the testifier gives the addressee a reason to believe—such as presenting oneself as having adequate epistemic reason to believe *p* or inviting the addressee to trust that *p*.<sup>18</sup> As Moran observes, assertions without any assurance occur in persuasion, argument, demonstration, therapy, examination, and interrogation (2006: 279–80; cf. Hinchman on conversational indirection, 2005: 568–72).<sup>19</sup> I would add that

222

the sentences we utter to teach young children language are usually assertions without assurance. It is compatible with these observations that assertion is an affirmation of a sort that normally expresses belief.<sup>20</sup> Various writers have insisted that assertion involves more than this—that it involves representing oneself as believing or as knowing.<sup>21</sup> Moran's examples of persuasion, argument, therapy, examination, and interrogation show that the 'knowing' variant is too strong, and examination seems to show that the 'believing' variant is also too strong. These claims very likely arise from confusing assertion with assurance. Nevertheless, they are compatible with both the assurance view and the assertion view I will eventually contrast it with, so we need not contest them here.

That assurance is more demanding than assertion is confirmed by the fact that the two are relationally different. Assuring entails asserting something *to* an *actual* addressee. If I say

something intending to assure Audrey, but she does not hear me, I do not assure her of anything—I do not make any assurance at all. But I do assert something. Asserting, then, does not require asserting to anyone. Indeed, it seems that it does not even require an *intended* audience. I can utter ‘I’m alone’ in the privacy of my house, and thereby assert that I’m alone. Even if the only possible point of such an utterance is to fix my belief that I am alone, it does not follow that I am addressing my assertion *to myself*. It seems unlikely that self-address is necessary or helpful for belief-fixation; verbal repetition does the job.<sup>22</sup>

### IS ASSURING ANALOGOUS TO PROMISING?

Having formulated the assurance view and remarked on the difference between assurance and assertion, we may now consider arguments for it.

Moran analogizes assuring in testimony to promising.<sup>23</sup> I take it that the analogy is supposed to lend significant support to the assurance view. The

223

proposed analogy is this. Suppose I promise you that I will weed the garden. Then my promise not only gives me a practical reason to perform the action of weeding the garden, but gives you a practical reason to rely on my performing this action when you plan your own actions. In promising, I make a commitment to perform the action, and you receive an entitlement to rely on my performing the action.<sup>24</sup> Analogously, my assuring you that it rained in Portland gives you an epistemic reason to believe this proposition. In assuring, I make a commitment that the proposition is true, and you receive an entitlement to believe it, at least if you accept my assurance. This is just what the assurance view claims.

I don’t deny that we can describe promising and typical cases of assuring in a parallel way, but the differences between these stand in the way of analogical support for the assurance view.<sup>25</sup>

First, a promise is made only if the promisee *accepts* the promise, hence only if there is an agreement to perform the promised action.<sup>26</sup> The promisee’s reason to rely on the promised action is no more given by the promise than by the acceptance—it is given by the agreement. The reason does not arise merely because in virtue of its content the promise offers a benefit that gives the recipient a reason to accept the promise and fulfil its terms. If I say ‘I promise to weed the garden’ to Elmer and Alma, and Elmer accepts my promise but Alma shrugs, Elmer will, and Alma will not, be entitled to rely on my performance of the action. If Alma rebuked me for not weeding the garden, I could respond, ‘But you just shrugged when I offered. So I never gave you any reason to rely on my weeding the garden.’ By contrast, in the case of assuring in testifying, I can assure you whether or not you accept my assurance. All that is required on your part is that you hear and understand what I say and recognize that I am assuring you. More importantly, if my assurance or its acceptance is to give you an epistemic reason to believe what I say, then my assurance all by itself must do so, regardless of whether you accept it. For when

I assure you that it rained in Portland, I do so before you accept my assurance. In assuring, I guarantee that the proposition is true. I have already stuck my neck out, and once I have done this, it is too late for me to pull it back in and deprive you of a reason to believe, if you should

224

decide not to accept my assurance. The guarantee all by itself gives you a reason to believe, if anything does. I will return to this point below.

This disanalogy between assuring and promising fits the offer-of-benefit agreement view, that testimonial reason-giving does not require an actual agreement but only a reason to accept the assurance because in virtue of its content it offers the addressee enough benefit. But of course that there is such a disanalogy reduces or undermines support for the *offer-of-benefit agreement* view from the comparison of assuring and promising. And obviously the disanalogy contradicts the *actual agreement* view.

There is a second disanalogy between assuring and promising, one that reduces or undermines support for Moran's *two-tiered assurance* view from the comparison of assuring and promising. My testimonial assurance, and your acceptance or reason to accept, are supposed to give you an epistemic reason to believe without giving you a *good* epistemic reason to believe. More is required for a good epistemic reason than assurance or its acceptance, as Moran allows. I must be sincere, responsible, knowledgeable, trustworthy, or reliable in assertion if a good epistemic reason is to be given at all.<sup>27</sup> Suppose I assure you about a subject, yet later it turns out that I was not responsible, or knowledgeable, etc. in asserting *p*. Then my assurance hasn't given you a good epistemic reason to believe *p*. Or so Moran allows in order to respond to the worry that assurance or its acceptance is not enough for a good reason. But my promise and your acceptance of it give you not only a (prima facie) practical reason to rely on the promised action, but a *good* practical reason to do so, whether or not I am responsible, and so forth, in promising. If later it turns out that I was not responsible, you still have good reason to rely on my promise; I am still bound by it, and you may hold me to it. The objection to the analogy, then, is that my promise and your acceptance by themselves give you more than my testimonial assurance and your acceptance do according to the two-tiered assurance view. If assuring were fully analogous to promising, the analogy would support a single-tiered view. Luckily for Moran's two-tiered view, the full analogy does not hold. Assurance or its acceptance depends on background conditions to give a good epistemic reason, and promising does not depend on such conditions to give a good practical reason. This disanalogy reduces or undermines support for Moran's two-tiered view from the comparison of assuring and promising.

This second disanalogy is related to a third. While testifying and promising both exhibit a variation from one instance to the next in the degree of assurance offered, the two differ in the way that the degree of assurance offered affects the quality of the reasons given. In the case of testifying, the greater the degree of assurance the better the epistemic reason

given, other things equal. But greater assurance in making a promise doesn't increase the quality of the

225

practical reason donated by the promise. The most tepid promise commits the promisor as firmly as the most emphatic promise. No doubt greater emphasis increases the quality of the promisee's evidentially based epistemic reason to *expect* the promised action to occur, and that may in turn increase the quality of the promisee's practical reason to rely on the promised action. But this increase in the quality of the practical reason to expect the action derives not from the promise and its acceptance but from whatever implicit testimony or evidence that the promised action will occur that the promisor incidentally gives while promising. The quality of the practical reason to rely on the promised action depends entirely on whether there is an agreement, not on the degree of assurance.<sup>28</sup>

These three disanalogies reduce or undermine altogether any support for Moran's version of the assurance view from the comparison of assuring and promising.

#### DOES AN ACTUAL AGREEMENT GIVE AN EPISTEMIC REASON?

Let us set aside the question of an analogy with promising and ask directly whether anything in assurance or its acceptance might generate an epistemic reason to believe the testimonial proposition. We will need to treat separately the two variants of the agreement assurance view, since they identify different sources of an epistemic reason. I will focus in this section on the *actual* agreement variant:

An actual agreement gives the addressee an epistemic reason to believe *p* because the assurance and acceptance together generate commitments for both the testifier and the addressee, and these commitments give the addressee an epistemic reason to believe.<sup>29</sup>

226

We are to consider, then, whether an *actual* agreement gives the addressee an *epistemic* reason to believe the testimonial proposition.<sup>30</sup> We may initially focus on just one proposal for what is involved in assurance—my inviting you to trust me that *p*—and then generalize what we find about this proposal. Does my inviting you to trust me together with your acceptance of my invitation give you an epistemic reason to believe *p*? We may grant that if you accept my invitation, then you have a (prima facie) reason to do what I invite you to do (in this case, trust me that *p*), just as one always has a reason to do what one is invited to do when one accepts an invitation. You have such a reason even if you lacked any reason to accept my invitation. By accepting my invitation to trust me, you commit yourself to trusting me and you have a reason to do so in virtue of this commitment. On the assumption that if something gives you a reason to trust me that *p* of the sort just mentioned, then it also gives you an epistemic reason to believe me that *p*, your acceptance of my invitation gives you an epistemic reason to believe me that *p*.

The trouble with this argument is that we cannot infer from the fact that your accepting my invitation to trust me gives you a reason to trust me, that it also gives you an



*epistemic* reason to believe me that *p*. The reason your accepting my invitation gives you is not itself an epistemic reason to believe *p*; and there is no basis for thinking that giving you a reason that is not itself epistemic could give you an epistemic reason. To show that the reason your accepting my invitation gives you is not itself an epistemic reason to believe *p*, one might observe that it does not bear on whether *p* is true, as required for an epistemic reason to believe *p*. Although I find this observation sufficient to make the point, an assurance theorist might reply by denying that an epistemic reason must bear on the truth, conceding only that an *evidential* epistemic reason must do so. Not knowing how to riposte a blank denial of the basic principle that an epistemic reason must bear on the truth, I support the point in a different way.

If the reason your accepting my invitation gives you is simply the sort you generally acquire from accepting an invitation, then it is defeasible by countervailing non-epistemic reasons. For example, it may be overridden by its being prudentially costly or risky for you to trust me—just as the reason you have to join me for dinner, given by accepting my invitation to dinner, may be overridden by its being costly or risky for you to join me. Suppose that after accepting my invitation to trust me, someone threatens to do you physical harm if you trust me. If the threat is credible and severe, the reason you have to trust me is defeated by prudential reasons you have not to trust me. But no epistemic

227

reason to believe me is defeated by such countervailing prudential reasons. So the reason your accepting my invitation gives you to trust me is not an epistemic reason to believe me. And there is no reason to think that your being given a reason to trust me gives you an epistemic reason to believe me. This is enough to cast doubt on the argument for the actual agreement view (or the version of it that traces the epistemic reason to an accepted invitation to trust).<sup>31</sup> In the absence of a persuasive argument for the view, we have no cause to take the view seriously.

I do not deny that the practical reason to trust given by an accepted invitation is a reason to do something—to trust me that *p*—that *facilitates* having an epistemic reason to believe *p*. For testimony gives an epistemic reason to believe *p* only if one takes up testimony in the right way, and trusting is one of the right ways to take up testimony. But this is only an instrumental relation between accepted invitations to trust and being given an epistemic reason.

I have considered the version of the actual agreement view that traces the epistemic reason to an accepted invitation to trust. But what we have said against the argument for that version generalizes to parallel arguments for other versions, including those that trace the epistemic reason to a presentation of an epistemic reason and to an offer to assume epistemic responsibility for asserting *p*. The reason given by an agreement in testimony is always defeasible by countervailing practical reasons. This is a consequence of the general fact that the reason given to fulfil any agreement whatever is always defeasible by countervailing practical reasons. No such reason is an epistemic reason to believe. And giving such a reason never gives an epistemic reason to believe.<sup>32</sup>

## DOES AN OFFER OF BENEFIT GIVE AN EPISTEMIC REASON?

I have rebutted the argument for the actual agreement view by contesting whether an actual agreement accounts for giving an epistemic reason. I turn now from the actual agreement to the *offer-of-benefit* agreement variant of the agreement view:

Assurance gives the addressee an epistemic reason to believe  $p$  because in virtue of its content it offers the addressee a benefit great enough to give the addressee a reason to accept the assurance and fulfill its terms.

On this view, the character of the benefit must make the reason to believe epistemic.

I see more hope for this variant than for the actual agreement variant. The hope is that, even though no agreement gives a reason not defeasible by countervailing practical reasons, an offer of a benefit may nevertheless do so, and an offer of the right kind of benefit may make this an epistemic reason to believe. The offer of the benefit must by itself provide a reason to believe not defeasible by practical considerations and aptly called epistemic. Now, there is no plausibility to the suggestion that my invitation to trust is such an offer. Any reason given by an invitation to do something would seem to be defeasible by countervailing practical considerations. But there is another candidate for what I offer in assurance that stands a better chance of providing the desired sort of reason. I have in mind my offer of my taking or assuming epistemic responsibility for asserting  $p$ . This is an offer to satisfy norms of epistemically responsible behaviour in asserting  $p$ . This offer would seem to be closer to a guarantee that  $p$  than the other components of assurance mentioned by Moran.<sup>33</sup> My offer specifies a benefit to you, of my taking epistemic responsibility for asserting  $p$ . The prospect of receiving the benefit of my assuming epistemic responsibility for asserting  $p$  conditional on your accepting my offer is supposed to entice you to accept my offer. And the benefit offered gives you a reason to believe  $p$ , since doing so is part of accepting the offer.<sup>34</sup> Since the benefit offered is epistemic (namely, my assuming epistemic responsibility), the offer might plausibly be taken to give you an *epistemic* reason to believe  $p$ . It is not obvious that your reason to accept the

offer can be defeated by countervailing practical reasons, as the reason given by my invitation to trust (and your acceptance of that invitation) is. In any event, I set aside the objection from defeasibility that afflicts the invitation proposal and press a different objection to the present proposal.

I read Moran as making this proposal in the only passage in his article that suggests why he takes assurance to give an epistemic reason:

If it seems difficult to see how anything, even someone's words, could acquire some epistemic value through something like conferral, perhaps because this suggests something too arbitrary or ceremonial to constitute a genuine reason for belief, it should

be remembered that for both parties this conferral is by its nature an overt assumption of specific responsibility on the part of the speaker. This is no more (or less) mysterious than how an explicit agreement or contract alters one's responsibilities, actions which are also within the capacities of ordinary speakers. (Moran 2006: 288–9)

On the most charitable reading of this passage, Moran proposes that my offer of taking epistemic responsibility for asserting  $p$  gives you an epistemic reason to believe  $p$  because it gives you a reason to accept the offer.<sup>35</sup> The reason it gives you to accept the offer is that doing so will bring the benefit of my taking such responsibility. And having such a reason gives you an epistemic reason to believe because the benefit offered is epistemic.<sup>36</sup>

Is this proposal correct? The final sentence of the passage most recently quoted from Moran suggests that all the assurance view need do to support an affirmative answer is to explain how your acceptance of my offer would alter my epistemic responsibilities. But to support the claim that the offer gives an epistemic reason, it is not enough to explain how my epistemic responsibilities would differ; it must be shown that the offer of such a difference gives you an epistemic reason to believe  $p$ . It must be shown that the offer of the benefit to you of my taking epistemic responsibility for asserting  $p$  gives a reason that, when good, contributes to your epistemic justification for believing  $p$ .<sup>37</sup>

## 230

The objection to the proposal I want to press is that the offer it attributes to me makes no sense.<sup>38</sup> Clearly it is logically impossible for me to offer you that the assertion I now make—one completed at the moment I complete this very offer—is epistemically responsible. For by the time I complete the assertion and the offer, it is too late for the epistemic responsibility of the assertion to depend on whether you accept the offer, as it must if my offer is to make sense. Rather, my offer must be that I will take epistemic responsibility in the future for my present assertion that  $p$ , where this means taking the consequences should it later be discovered that my assertion was not epistemically responsible at the time I made it. (Presumably, taking the consequences amounts to admitting that my assertion was not epistemically responsible and making reparations for any damage from your trusting that  $p$  under those conditions.) Now, it is logically possible for me to make this offer. However, you can have no reason to accept it if your aim is merely to acquire an epistemic reason to believe  $p$ . This offer gives you no more epistemic reason to believe  $p$  than would my offer to flip a coin and in exchange for your trusting that  $p$  if it lands Heads and not- $p$  if it lands Tails, take the consequences should it later be discovered that the proposition in which you trust is false.

Granted, if you had epistemic reason to believe that I would not offer to take the consequences should it later be discovered that my assertion was not epistemically responsible *unless* my assertion *were* epistemically responsible, this would give you an evidential reason to believe that my assertion was epistemically responsible, and this in turn might very well give you an evidential reason to believe  $p$ . But Moran proposes that my offer gives you a *non-evidential* epistemic reason to believe  $p$ . And my merely

offering to take the consequences does not by itself give you any such reason. Nor is it plausible that such an offer provides an epistemic reason to believe *p*, though if it were later discovered that my assertion was not epistemically responsible, it would be a *bad* epistemic reason. For whether the reason given at the time of testimony is good or bad should not depend on whether it is later discovered that my assertion was not epistemically responsible. The key point, however, is that my offer is missing something crucial for giving a reason. One might say that what is missing is a guarantee that my assertion was epistemically responsible at the time. But that simply names the problem Moran is trying to solve by specifying that the relevant offer is to take the future consequences of my

231

present assertion—an unsuccessful proposal, as I have argued. The offer Moran cites is not one that could give you an epistemic reason to believe *p*.<sup>39</sup>

In short, on no version of what is involved in assurance and its acceptance do either of these give you an epistemic reason to believe *p*, at most a practical reason to accept an invitation to trust or to accept an offer to take epistemic responsibility in asserting *p*. Neither the actual acceptance nor the offer-of-benefit assurance view is correct. Assurance and its acceptance provide no epistemic reason to believe that is not present in other cases of assertion.

It is fair for the assurance theorist to respond at this point by asking what benefit a practice of assurance can be thought to confer on our communicative-cognitive economy, if not the benefit of donating epistemic reasons? I would suggest that it confers benefits of two sorts.

First, as I have explained, an invitation to trust can bind the addressee who accepts it in the usual way that acceptances of invitations bind the invitee. The social desirability of accepting an invitation to trust may motivate the recipient to accept the invitation, and the binding then motivates the addressee to follow through and trust. This is beneficial to the extent that testifiers give assurances only of true beliefs. The fact that addressees tend to believe assurances motivates testifiers to be cautious in their assurances. To the extent that people desire to be able to give assurances, this desire motivates them to acquire true beliefs, and toward that end justified beliefs, in propositions for which they then provide assurances. Similarly, an offer to take future responsibility for what is assured can motivate addressees to accept the offer as insurance and thus to trust, and this in turn motivates testifiers to be cautious in their assurances and acquire true and justified beliefs.

Second, assurance itself (weakly) indicates the truth of the assured proposition. That someone is willing to offer an assumption of responsibility in exchange for trust indicates his or her confidence in the truth of the proposition, and confidence in the truth of a proposition indicates that it is true. In virtue of this, our recognition of assurance gives us

an epistemic reason to believe (whether an evidential reason or some other kind). These benefits of assurance do not depend on the truth of the assurance view.

#### IS ASSURANCE NECESSARY FOR GIVING A TESTIMONIAL REASON?

It is natural to include in the assurance view the position that assurance is necessary and not just sufficient (as we have taken the view) for giving a

232

non-evidential epistemic reason. And it is equally natural to object to this position on the ground that testimony gives an overhearer of the testimony the same epistemic reason to believe  $p$  as it gives the addressee. Suppose a testifier assures an addressee and as a consequence, whether foreseen or not, leads an overhearing bystander to believe  $p$ . The objection is that the testimony gives the overhearer the same epistemic reason to believe  $p$  as it gives the addressee. Yet the testifier does not assure the overhearer that  $p$ , present any reason to believe  $p$  to the overhearer, invite the overhearer to trust him, nor offer to take epistemic responsibility for the assertion, and the overhearer therefore cannot, and has no reason to, accept any invitation or offer. So testimony can give a hearer an epistemic reason to believe  $p$  even though the testifier does not assure the hearer, does not lend the hearer an opportunity to accept any assurance, and the hearer does not accept any assurance. I will press this objection to the assurance view.

Lackey sets out an overhearer case:<sup>40</sup>

Ben and Kate, thinking that they are alone in their office building, are having a discussion about the private lives of their co-workers. During the course of their conversation, Ben tells Kate that their boss is having an affair with the latest intern who has been hired by the company, Irene. Unbeknownst to them, however, Earl has been eavesdropping on their conversation and so he, like Kate, comes to believe solely on the basis of Ben's testimony—which is in fact both true and epistemically impeccable—that his boss is having an affair with Irene. Moreover, Kate and Earl not only have the same relevant background information about both Ben's reliability as a testifier and the proffered testimony, they also are properly functioning recipients of testimony who possess no relevant undefeated defeaters. (Lackey 2008: 241)

On the assurance view of testimony, Ben's assurance or Kate's acceptance of it gives Kate an epistemic reason to believe that the boss and Irene are having an affair, but it does not give Earl an epistemic reason to believe this. For Earl is given and accepts no assurance. Earl may have an evidential reason to believe the proposition, based on his belief that Ben assures Kate and that Ben is a trustworthy and reliable source of information of this sort when he makes assurances in circumstances of this kind. But Earl has no reason from assurance. Thus, Kate is given and has an epistemic reason to believe that Earl is not given and does not have. But this implication of the assurance view is counter-intuitive. Intuitively, Earl is given and has the same epistemic reason to believe that Kate does, even though Kate is assured and Earl is not.<sup>41</sup>

233

There are several ways the proponent of the assurance view of testimony might reply to this objection. One is to stonewall and insist that Kate has an epistemic reason that Earl lacks. According to Ross, for Kate not to believe Ben would be for her to ‘challenge’ Ben’s authority, but this would not be so for Earl.<sup>42</sup> Moreover, Kate would be entitled to resent Ben if it turned out he was not trustworthy on the matter, but Earl would not be so entitled. According to Hinchman, Ben would be entitled to feel slighted if Kate did not believe him, but not if Earl did not believe him. Ross and Hinchman take these points as evidence of a difference in the non-evidential reasons Kate and Earl are given. If Earl were given the same reasons by Ben’s testimony as Kate is, then his not believing Ben would challenge Ben’s authority, he would be entitled to resent Ben if it turned out that Ben was not trustworthy, and Ben would be entitled to feel slighted if Earl did not believe him. The power of challenge and the entitlements to resentment and feeling slighted are powers and entitlements that typically arise from agreements. So it is natural to propose that Kate’s agreement with Ben gives rise to these powers and entitlements and that Kate has a reason in virtue of that agreement that Earl does not have.

Lackey questions whether the cases of Kate and Earl really differ in these ways. If Earl were not to believe Ben, then, according to Lackey, he too would challenge Ben’s authority in the sense that his action would reveal his belief that Ben is untrustworthy. And Earl would also be entitled to resent Ben if it turned out Ben was not trustworthy. And if Earl were not to believe Ben, then Ben would be entitled to feel slighted.

Ross and Hinchman may respond persuasively to this criticism by allowing that all this is true, but maintaining that there is still a crucial difference between Kate and Earl: Kate could *specialy* challenge Ben’s authority and would have a special reason to resent Ben, but Earl could not, and Ben would be *specialy* entitled to feel slighted by Kate, but not by Earl. Earl’s challenge and reason to resent Ben, and Ben’s entitlement to resent Earl, fall under norms of assertion rather than norms generated by assurance and acceptance. If Earl were not to believe Ben, Earl would challenge Ben’s authority in the sense that not believing Ben manifests that Earl does not take the relevant norm of assertion (‘One should believe what people assert when one takes them to be authoritative’) to apply to this instance of assertion. Earl’s not believing Ben would manifest that Earl does not believe Ben to be authoritative. And Earl would have reason to resent Ben if it turned out that Ben is untrustworthy because then Ben would violate a different norm of assertion (‘One should not assert something about which one is not trustworthy’). Earl would have the sort of reason one has to resent the behaviour of those who violate norms that do not arise from agreements. Finally, if Ben were to discover that Earl does believe him, then Ben would be entitled to

feel slighted because by not believing him Earl manifests that he does not believe Ben to be authoritative.

But, Ross and Hinchman may legitimately insist, these reasons and entitlements deriving from norms of assertion are quite different from the reasons and entitlements that arise

from Ben's assurance and Kate's acceptance. In assuring Kate, Ben makes an offer for Kate to rely on his word by trust alone and an offer to fulfil the obligation imposed by making this offer only if he is trustworthy. If Kate were not to believe Ben because she does not accept Ben's assurance, then she would challenge Ben's authority, not only in the sense in which Earl would challenge it were he not to believe, but in an additional sense: not to believe Ben would indicate not accepting Ben's assurance, and this would manifest that Kate does not believe Ben to be authoritative, since it manifests that she does not take the relevant norm of assurance ('One should believe what people assure when one takes them to be authoritative') to apply to this instance of assertion. And if Kate were not to believe Ben after accepting his assurance, then she would challenge Ben's authority in the different sense that she manifests that she regards the agreement they have made to be illegitimate, presumably because she does not take Ben to be trustworthy. And if it turned out that Ben is not trustworthy, Kate would be entitled to resent Ben, not merely for violating the norm of assertion to which we referred above, but for reneging on his accepted offer. And finally, if Kate were not to believe Ben, then Ben would be entitled to feel slighted because Kate reveals her belief that Ben is not trustworthy, either by not believing him once she has accepted his offer, or by refusing to accept his offer. These appear to be genuine differences between Kate and Earl. Moreover, these differences arise from the fact that Kate accepts Ben's assurance, while Ben does not give his assurance to Earl, and Earl accepts no assurance.

But this response to Lackey shows only that there are differences of these sorts between Kate and Earl, and that the differences stem from differences in assurance and acceptance. There remains an objection to the assurance view, also made by Lackey, that these differences are not *epistemically* relevant; they do not show that Ben's testimony gives a non-evidential epistemic reason to Kate but not to Earl. I have already supported the same conclusion in the preceding sections. Kate has a practical reason to trust that Earl lacks. But this difference between them does not entail that Kate has an epistemic reason to believe that Earl lacks.

Moran defends the assurance view from the objection from overhearers in a different way:

while the overhearer may get a reason to believe without having the right to complaint that is conferred on the addressee, the fact that the overhearer of the assertion acquires any reason to believe from listening to these words is dependent on them being addressed to someone, with the force of assuming responsibility and thereby conferring a right of complaint . . . Without that [dependence], the question of what speech act, if any, is

being performed with these words would not be settled, and hence the overhearer could not get started on assessing their epistemic significance. (Imagine overhearing someone say 'The rain in Spain falls mainly on the plain'. Until you know what speech act, if any is being performed here, you don't know if considerations of reliability or trustworthiness are even relevant to the status of the words as source of knowledge about the weather in

Spain.) So, while in both cases (promising and telling), the overhearer can gain a reason to believe something without entering into the normative relation of promisor-promisee or teller-believer, in the overhearing of testimony he only gains a reason to believe something because such a relationship *has* been established by the original speaker and addressee.

(Moran 2006: 296)

Here Moran maintains that the overhearer gains a non-evidential epistemic reason to believe *p* in virtue of knowing that the testifier assures and the addressee accepts this assurance—in virtue of the normative, or personal, relation between the testifier and the addressee—despite not receiving or accepting that assurance. Moran's analogy to the utterance about Spain suggests that his argument for his position is that the overhearer wouldn't gain an epistemic reason to believe if he or she didn't know that the testifier assures and the addressee accepts this assurance.

A preliminary point about Moran's argument is that (in a typical case) to gain an epistemic reason, the overhearer need know at most that the testifier assures, not that the addressee accepts the assurance. This is indeed what Moran should say, given that (as I argued earlier) he should hold the offer-of-benefit agreement view, not the actual agreement view. On the offer-of-benefit view, all that is needed for the addressee to be given an epistemic reason to believe is the testifier's assurance. So all that is needed for the overhearer to have an epistemic reason to believe is knowing that the testifier assures. The overhearer need not know of a personal relationship between the testifier and the addressee. Indeed, there need be no such relationship. So far, there is no objection to the argument. The objection comes when we see that the premise of dependency does not at all entail the conclusion of a *non-evidential* epistemic reason. The overhearer's having an evidential epistemic reason would also depend on knowing that the testifier assures. So Moran cannot infer from a dependency to a non-evidential epistemic reason.

Moran would have to reply that the overhearer must have a non-evidential epistemic reason because the epistemic reason the overhearer has depends on knowing that there is assurance and not merely assertion. But, Moran must insist, if the overhearer merely has an evidential epistemic reason, having this reason could just as well depend on knowing that there is assertion as on knowing that there is assurance. The fact that it does depend on knowing that there is assurance can be explained only by supposing that the epistemic reason gained is non-evidential.

The claim is that a dependency on knowing that there is assurance can be explained by supposing the reason is non-evidential, while a dependency on knowing that there is assertion cannot be so explained. But can a dependency on knowing that there is assurance be explained by supposing the reason is non-evidential? Suppose that (contrary to my earlier argument) the assurance theorist is right in claiming that we can explain how assurance gives the addressee a non-evidential epistemic reason. The explanation is



that assurance involves an offer to take epistemic responsibility that entitles the addressee to rely on its being the case that  $p$ . But this explanation does not transfer to the overhearer. The offer is not made to the overhearer, so confers no title in the overhearer to rely on the testifier's taking epistemic responsibility for asserting  $p$ . So we have no explanation of how assurance provides the overhearer with an epistemic reason, no explanation that entails that the reason is non-evidential. Hence, we have no ground to claim that a dependency on knowing that there is assurance can be explained by supposing that the reason is non-evidential. Similarly, we have no ground to claim that a dependency on knowing that there is assertion cannot be so explained. For as far as we can now tell, when we finally obtain an explanation of how assurance provides the overhearer with an epistemic reason, we may convert it to a parallel explanation of how assertion provides the overhearer with an epistemic reason, and that explanation does not suppose that the reason is non-evidential.

Perhaps the more important objection to the envisioned reply is that the overhearer's having an epistemic reason does not depend on knowing that there is assurance as opposed to mere assertion. It seems that an overhearer of persuasion, examination, interrogation, and the like can acquire a reason to believe that does not differ significantly from one acquired from assurance. Suppose you overhear someone answer another's question, and you assume, plausibly enough, that the answerer assures the questioner. You may well have an epistemic reason to believe the answer. Suppose you later find out that the person was not assuring that  $p$  but rather answering an exam question. This news would not by itself lead you to concede, 'Well, then, I must not have come to have an epistemic reason from what they said'. Indeed, it does not seem crucial for receiving an epistemic reason to believe that you even have considered whether the case was one of assurance or mere assertion in an examination. Of course, if you discovered that the examinee was a beginning student of the subject and the question overheard was the most difficult question on the test, you would reduce your confidence in their answer and judge that your reason isn't as strong as you thought, though I don't think you would necessarily have to retract your judgement that you acquired an epistemic reason. A proponent of the assurance view might protest that what you receive in such cases is merely an evidential reason to believe. But where is the argument that this is any more true in these cases than in that of the overhearer of assurance?

#### CAN THE ASSURANCE VIEW ACCOUNT FOR ALL NON-EVIDENTIAL, NON-REDUCTIVE TESTIMONIAL REASONS?

A final reservation about the assurance view is that it is incomplete given one of its motivations: it cannot account for all non-reductive epistemic testimonial reasons if the regress objection to reductive views of testimonial justification is successful. According to the regress objection, reductive views are mistaken because we do not have enough experience, or enough non-testimonially justified beliefs, on which to base all of our testimonially justified beliefs. Some testimonially justified beliefs are therefore not based

on these non-testimonial sources of justification. The assurance view cannot explain the justification of all of these non-reductively testimonially justified beliefs. For the acceptance of assurance requires the recognition of assurance—or of what is involved in it, the testifier's presentation of a reason, invitation to trust, or offer to assume epistemic responsibility. So on the actual acceptance view, giving an epistemic reason to believe requires actual recognition of assurance or the involved actions. And on the offer-of-benefit view, it requires that the addressee is assumed by the testifier to be able to recognize assurance. But according to the regress objection, we do not have enough non-testimonially justified beliefs about intentions to afford a basis for our recognition of testimonial intentions in all instances of testimonially justified belief. Thus, if the regress objection establishes that there is some non-reductive testimonial justification, it also establishes that we do not recognize assurance in all instances in which we have testimonially justified beliefs. So it establishes that the assurance view gives an account of only some instances of testimonially justified belief. The assurance view does not apply to all *primitive instances* of testimonially justified beliefs (where primitive instances are those in which the hearer does not already have enough non-testimonially justified beliefs to recognize assurance). Thus, an assurance view motivated in part by the regress argument must be supplemented with another view of testimonial justification to cover all primitive instances.

The proponent of the assurance view might respond that there must be something wrong with this objection to the assurance view. For on any tenable view of testimonial justification, such justification must involve the testifier's assertion that *p* and the hearer's recognition that the testifier asserts *p*. Yet there is just as much reason to doubt that there is a non-testimonial basis for our recognition of assertion as for our recognition of assurance: each requires recognition of speaker intentions and applicable norms; you recognize that the testifier asserts *p* only if you recognize that the testifier has certain communicative intentions and is governed by norms of assertion. But then the premises of the

238

consideration against the assurance view lead to scepticism about testimonial justification.

In reply, I find it plausible that on any tenable view testimonial justification must involve the testifier's assertion that *p*, but I doubt that it must also involve the hearer's recognition that the testifier asserts *p*. It is enough for the recipient to believe *p* because the testifier asserts *p*, in a manner analogous to that in which remembering *p* requires only believing *p* because one has believed *p*. Just as one can remember *p* without recognizing that one once believed *p*, so one can believe on testimony without recognizing that the testifier asserted *p*. This is borne out by the fact that much of what we believe we absorb 'from the air', without noticing that anyone has asserted the relevant propositions. We can only surmise that we hold these beliefs because others asserted these propositions. But we are

testimonially justified in these beliefs, regardless of whether we recognize that anyone asserted them.

I don't think the assurance theorist can respond to this by proposing that on the assurance view the hearer need only believe  $p$  because the testifier assures  $p$ , in a manner analogous to remembering  $p$ , or knowing  $p$  now because one knew  $p$ . The story the assurance view tells about how the testifier's presentation (or the like) gives the hearer a reason is implausible unless it is assumed that the recipient recognizes that there is a presentation. No reason to believe on the model of the assurance view is given if I utter ' $p$ , I assure you', but you are not in a position to accept my assurance because, for want of non-testimonial justification, you are not in a position to recognize that I am making a presentation. I conclude that in this regard the assurance view labours under theoretical limitations that do not encumber other tenable non-reductive views of testimonial justification. These views can aim at full generality where the assurance view cannot.

#### IF NOT ASSURANCE, THEN WHAT?

Our discussion suggests the following condition on testimonial reason-giving: testimony that  $p$  gives the hearer an epistemic reason to believe  $p$  only if the hearer believes  $p$  because the testifier asserts  $p$ . Call this the *assertion view* of testimony.<sup>43</sup>

239

On this view, testimony gives such a reason only in cases of assertion. These include not only cases of assurance, but of persuasion, examination, linguistic training, and so on.<sup>44</sup> A reason is given only when the testimony is taken up in a certain way. A major burden for the assertion view is just what sort of uptake is required for giving a reason. It would seem not to be enough for the recipient merely to hear the testimony. But it is too strong to require that the hearer recognizes that the testifier asserts  $p$ , as 'in the air' cases and primitive instances of testimonial justification show. There are several possible sources of inspiration in looking for the right requirement. One is an analogy with what is required for memorial reason-giving. Memory gives me an epistemic reason to believe only if I believe because I believed, but the former belief must arise from the latter in a proper way. Another possible source of inspiration is that a substantive condition on justification like reliability may limit the sorts of uptake that can give rise to testimonially justified belief and in this way render unnecessary any further constraint to rule out exceptionable cases. But I will not pursue here the difficult question of what sort of uptake is required.<sup>45</sup>

The assertion theorist can allow that, where there is assurance, the degree of assurance is a factor in the quality of the reason given—that, other things equal, believing  $p$  from assurance gives the recipient a better epistemic reason to believe  $p$  the greater the testifier's assurance. The quality of the epistemic reason would be a function of all of the good-making properties. Obviously, it is a major project to define this function. But the assurance theorist who, like Moran, wishes to concede that properties other than assurance or acceptance are good-making is burdened with the same project.

240

It is natural to think that a reason must have representational content of some kind and that the content of a testimonially given reason must be one the hearer is in a position to recognize (i.e. recognize as being the content of the given reason). The assurance view has a ready answer to what the content is—namely, ‘the testifier assures that  $p$ ’. It also requires that the addressee is in a position to recognize that the testifier assures that  $p$ . This requirement takes us close, though not all the way, to the requirement that the addressee is in a position to recognize the content as being the content of the reason given. One might take this approximation to the latter requirement to be an advantage of the assurance view over the assertion view. But in fact any view that endeavours, as the assertion view does, to cover primitive instances of testimony will have to swallow the idea that there are instances of testimonial reason-giving in which the hearer is not in a position to recognize what content the given reason has or that it has the content of the given reason. One candidate for the content of a testimonial reason is ‘the testifier asserts  $p$ ’. Another, less plausible candidate is the content of some reason the testifier possesses that favours  $p$ . But hearers are not in a position to recognize these in all instances of testimonial reason-giving. The assertion view may well have to deny either that a reason must have representational content or that the content of the reason must be one the hearer is in a position to recognize. This is a cost of any account of testimonial reason-giving that covers the primitive instances.<sup>46</sup>

There is a lingering question: if assurance isn't required to give the hearer an epistemic reason to believe, why should even assertion be required? Lackey (2008) argues persuasively that insincere assertion (in which the testifier does not believe  $p$ ) is enough to give testimonial knowledge. But if insincere assertion is enough to give a reason, why shouldn't any saying that  $p$ , whether an assertion or not, also be enough? A natural answer is that what counts as a vehicle for giving an epistemic reason is the sort of thing that generally gives a good reason, and what generally gives a good reason is assertion. Requiring the vehicle to be a sincere assertion is too strong because it happens often that testifiers possess good reasons to believe what they (for social or practical reasons) assert, even though for epistemically irrelevant reasons they do not believe those assertions. Allowing the vehicle to be less than an assertion is too weak because it happens very often that testifiers lack good reasons to believe what they merely say. Whether a saying is an assertion is a fairly accessible indicator of the presence of good reasons.

None of this is to deny that assurance and its acceptance play a role in the testimonial creation of reasons beyond what assertion plays. They generate practical reasons to trust given the aim of trusting what is truly and responsibly

241

asserted. The generation and transmission of these reasons motivate both the testifier and the hearer to a higher level of responsibility in assertion, testimonial uptake, and follow-up than they would otherwise exhibit. In these ways, assurance and its acceptance contribute to the acquisition of beliefs for which testifiers have a good epistemic reason.<sup>47</sup>

242

## 12 The Epistemology of Silence

Sanford C. Goldberg

1.

Consider the following two sorts of case.

CASE 1. In the course of a conversation, the issue comes up whether Jones regularly wore a bright red suit. You think about it come to the belief that, no, he did not. You reason that if he had regularly worn such a suit you would have seen him in it, and you can recall nothing of the sort.

CASE 2. In the course of a conversation, the issue comes up whether weapons of mass destruction were ever found in Iraq after the second Gulf War. You think about it and come to the belief that no such weapons have been found to date. You reason that if they had been found you would have heard about it by now, and you have heard nothing of the sort.

In this chapter I aim to compare and contrast these two sorts of belief based on the ‘silence’ of a trusted source.

Let us begin with the most salient features of the comparison. In both cases two salient points emerge immediately. First, the subject does not have any relevant confirming memory (of having seen Jones wearing such a suit; of having been informed that WMDs were found in Iraq). And second, the subject assumes that if the relevant proposition (that Jones regularly wore the suit; that WMDs were found in Iraq) were true, she would have obtained the relevant information (through perception or testimony) and would have retained that information in memory.

There are other, perhaps less salient similarities between the cases. For one, in both cases the proposition believed is a negation: Jones *didn't* regularly wear such a suit, WMDs have *not* been found to date. For another, in both cases the corresponding affirmative proposition is one regarding whose truth you might well have assigned (or have been disposed to assign) some low probability *prior to* reflecting on whether you had seen or heard about such a thing. For example, if Jones is not known for sartorial flamboyance, the hypothesis that he regularly wore a bright red suit would get a low prior probability (even before you reflected

on whether you had ever seen him in one, etc.). Or if you are deeply suspicious about Bush's motives in starting the second Iraq war, you might well have doubted the existence of WMDs in Iraq, even before reflecting on whether you had ever come across a report of such. In this light, the fact of ‘silence’ in the sources that are relied upon

should be seen as bearing on a proposition regarding whose truth you already have other (indirect) evidence.

So far we have been attending to the similarities between the cases, but there are differences too. The most obvious one concerns the identity of the information source being relied upon. In CASE 1 you are relying upon yourself, whereas in CASE 2 you are relying in addition on someone (or ones) other than yourself. In the latter case your ability to identify the relevant source or sources may be quite limited—you may only know that *some source or other* would have found out about the discovery, and would have reported it—or it may be relatively informed—you may know the identity of the CNN reporter on whom you rely for news from the Middle East, and in fact may know quite a bit about the CNN research team as a whole. But this difference between CASE 1 and CASE 2 may not be that significant. Granted that it is only in situations like CASE 2 that the question can meaningfully arise regarding the identity of the source on whom you are relying, even so, in many cases of this type there may be no question in your mind regarding whom you are relying upon (you know you are relying upon the reporters of the *New York Times* for your news). What is more, while it is true that in CASE 2 there can be variability in how much the subject might know about (the reliability<sup>1</sup> of) the relied-upon source (you might fail to know much about the *Times*' methods of gathering and disseminating that information), this might well be a similarity with CASE 1, as you may fail to know very much about how you gathered and stored the relevant information about Jones' clothes. To be sure, you will probably know enough in CASE 1 to say that your methods involve perception and memory; but at this level of generality the cases might well remain similar, since you might say in CASE 2 that the *New York Times* gathers its information by placing journalists in the field, and then disseminates the information through the publication of the daily paper (in hard copy or on the web). This leaves open precisely how much you know about (the reliability of) these processes you are designating as 'perception' and 'memory', 'journalistic practices', and 'publication of the daily paper'.

How should an epistemic assessment of these beliefs accommodate the similarities and differences between the individualistic case (CASE 1) and the social case (CASE 2)? In what follows I will argue that in both cases the beliefs themselves involve the same potential sources of unreliability,

244

and depend for their reliability on the goodness of the 'coverage' that is provided.

2.

My ultimate aim in this chapter is to characterize the epistemic significance of the 'silence' of a relied-upon source, whether that source is oneself (CASE 1) or another/others (CASE 2). To appreciate the epistemic significance of silence, however, it will help to begin with the notion of *epistemic reliance on an information source*, as this bears on cases in which an information source, far from being silent, offers a 'report'.

Once again we can distinguish the case of self-reliance—for example, the case of having a memory impression as of it being the case that *p*—and cases in which the epistemic reliance is on another subject as information source—the familiar case of testimony. After examining these cases, we will then return to the cases where the source in question is silent.

## 2.1.

Take the self-reliance (memory) case first. What sense does it make to say that in cases of memory one relies on oneself *qua* information source? Confronted with an apparent recollection whose content is *p*, one is relying on oneself in two ways. One is relying on oneself, first, as *the source of the information* that *p*: here one relies on the reliability of the processes through which one originally *acquired* the belief that *p*. In addition, one is also relying on oneself as responsible for *the reliable transmission* of a previously acquired belief-content: here one relies on one's memory to have succeeded in transmitting—retrieving and making available for one's present use—the very content one previously acquired, such that one wouldn't currently have the memory impression as of *p* unless one acquired that very information, *p*, at some earlier time. Hence we can distinguish reliance on oneself *qua* source of information,<sup>2</sup> from reliance on oneself *qua* transmitter (or preserver) of information.

245

Parallel distinctions can be made in the interpersonal case, where one's reliance is on an information source other than oneself. Here too we can distinguish two dimensions of reliance: on the source, and on the process of content preservation (or content transmission).<sup>3</sup> With respect to the latter sort of reliance, the parallel rests on the idea that we can regard proffered testimony, just as we regard apparent recollection, as serving the role of transmitting or preserving previously acquired information. Just as it is clear that memory (*qua* would-be preservative process) can fail to preserve the content of a previously acquired belief, in a way that does not impugn the reliability of the belief on acquisition, so too it is clear that the process by which a journalist's report is publicized (*qua* would-be preservative process) can distort the report, in a way that need not reflect poorly on the reliability of the original report.

So far I have been arguing that there are two dimensions along which we can develop the parallel between the sort of epistemic reliance in play in memorial belief and that which is in play in testimonial belief. In particular, both involve (i) an original source of information (ii) whose contents are presently transmitted (made available) to the subject. This distinction between (i) and (ii) will bear a good deal of weight in what follows.

In the case of testimony the distinction is clear enough: one's source for the information in question is another person, and that person might or might not choose to provide her audience with the information in question, and might or might not be reliable in the transmission. In the memory case the distinction between (i) and (ii) may not be so clear, but it is present. Perhaps everyone will agree that self-reliance in memory cases includes

reliance on oneself for information preservation (or what Burge (1993) calls ‘content preservation’). In recollecting that *p*, the information that *p* is retrieved and (through the memory impression) made available for the subject's present purposes. But self-reliance in cases of memorially sustained belief goes beyond this, to include reliance on oneself *qua* source as well. That is, one is not merely depending on oneself to have retrieved information that one in fact had acquired earlier (one doesn't want one's memory just making things up); one is also depending on oneself, or on one's past self, to have reliably acquired that information right at the outset (beliefs acquired through wishful thinking, even when successfully preserved in memory, are none the more reliable for having been so preserved).<sup>4</sup>

I have been arguing that in both memorial belief and in testimonial belief, we can factor the epistemic reliance involved into reliance on (i) the source of information and (ii) the transmission process whereby information acquired at some earlier time is presently made available to the subject. In connection with

246

(ii), however, there would appear to be several differences between the memory case and the testimony case.

One such difference concerns the link between the message and its recipient. Clearly, the link between the two can be severed in testimony cases: the source(s) on whom one relies might well make a report without one's coming across that report.<sup>5</sup> In memory cases, however, the link would not appear to be severable: it makes little sense to say that one's memory system has issued an apparent recollection which one oneself has failed to ‘come across’. (One might try to make sense of this in terms of different parts of the mind/brain failing to communicate; but as this is a matter for empirical theories of memory, not for epistemology, I will not pursue this idea here.) However, this difference, though real, need not impugn the parallel I have been developing between memory and testimony. Even if it is true that there are no ‘memory transmissions’ that fail to be apprehended by the subject herself, whereas there are ‘testimonial transmissions’ that can fail to be apprehended by her (as when she is not near the testifying source at the time of the testimony), still the distinction between source and transmitter remains in the memory case. This is seen in a case in which one oneself has information stored in memory, yet fails to be able to access it at some later time. Such cases are naturally thought to be the analogue of the case in which a subject reliably believes that *p* but does not express this belief to her audience.

A second, related difference between the nature of information transmission in memory and testimony cases concerns intermediaries in the process of transmission. Whereas few of us get our news ‘straight from the horse's mouth’, and so stand in need of intermediaries in the chain of communication that links our source's testimony to us, there are no such intermediaries in the case of memory: nothing comes between you and your apparent recollections. Again, however, this difference does not impugn the relevance of the distinction between (i) and (ii) for memory cases and testimony cases



alike. Rather, it merely suggests that the process of transmission in testimony cases can be significantly more involved than it is in memory cases.

Finally, a third apparent difference along the dimension of information transmission concerns the need for message comprehension. Whereas testimony, if it is to be useful to a subject, must be comprehended by her, it is by no means obvious that there is any intra-personal analogue of comprehension in cases of apparent recollection. The issue is somewhat complicated. Suppose (for example) that at least some memories are stored via sentential (or sentence-like)

247

vehicles. And suppose that for a public-language speaking subject the sentential vehicle is not a sentence of mentalese but rather a sentence from one's public language—the sentence of the language one would use to express the content of the recalled belief. In that case we might think that these 'stored sentences' require comprehension just as much as do pieces of testimony. Indeed, Burge 1999 appears to suggest a variation on this theme:

Comprehending standing, conceptual aspects of one's own thought and idiolect is itself, as a matter of psychological and sociological fact, normally dependent on having comprehended thoughts (one's own) that were shaped and expressed by the words of others. (Burge 1999: 243)

(What is interesting here is Burge's use of 'comprehension' and its cognates, in connection with one's own thoughts.) The matter is vexed, however, and I would not want to put too much weight on this aspect of the alleged analogy between memory and testimony. But once again it is worth pointing out that even if memory is disanalogous from testimony on this score, in that memory does not involve any 'comprehension' of a message, this need not undermine the claim that the epistemic reliance involved in both memory and testimony can be factored into reliance on (i) and on (ii). All that it would suggest is that the process of transmission differs between the two cases.

In sum, even if all of these alleged differences between memory and testimony are conceded to be actual differences, such a concession would only establish that the process of content transmission is more complicated (and allows for more distinctive kinds of error possibility) in the case of testimony. These differences in question are not sufficiently great to call into question the parallel between memory and testimony, nor are they sufficiently great to undermine the claim that in both types of case epistemic reliance can be factored into (i) and (ii). What is more, we might even think that whatever differences there are between testimony and memory in the above respects are to be expected, given the differences between intra-personal and interpersonal processes of information transmission. For it would seem that all three of the recently-noted differences flow from the difference between a system's retrieval of information that is *already stored within the system itself*, and a system's reception of information proffered by some source *other than* the system itself. The latter sort of case involves the

possibilities of ‘missing’ the message altogether, of the utility of communication chains involving intermediaries, and of the need to ‘decipher’ and ‘re-encode’ the message once it is received. These differences would thus appear to point to two distinct species (intra-personal and interpersonal) of a single underlying genus (information transmission).

## 2.2.

So far we have been considering epistemic reliance in cases in which a subject receives a ‘transmission’ from an information source. But what should be said

248

of epistemic reliance in cases in which a relied-upon source is *silent*? These would appear to be importantly *unlike* the standard cases in the epistemology of memorial or testimonial belief: arguably, the *lack* of a memory impression, or of a report, should not itself be thought of as something that the source in question ‘presents’.

This initial difference would appear to make an epistemic difference. For while the fact that you have an apparent recollection to the effect that *p*, or that you have observed a report to the effect that *p*, might (together with whatever justifies you in accepting the ‘presentation-as-true’) ground a justification for your sustaining or acquiring the belief that *p*, it is not clear how the fact that you have *no* apparent recollection to the effect that *p*, or that you have *not* observed a report to the effect that *p*, might (together with other things) ground a justification for your acquiring a belief in  $\sim p$ . After all, a popular (though not universally endorsed) claim among people working on testimony is that, when successful, testimony affects a *transmission of epistemic status*, from the testimony, to the hearer (or the hearer's belief). Since no such transmission is affected in cases in which there is no testimony (or other sort of report) in the first place, it is not immediately clear how to account for the epistemology of silence.<sup>6</sup>

The issue before us, then, is this: how (if at all) does the epistemic status of a hearer's belief depend on the epistemic perspective of those on whom she is epistemically relying, in those cases in which her sources are silent? This issue can be helpfully framed by comparing two subjunctive conditionals, one that (I will suggest) lies at the heart of the epistemology of testimony (and memory), and the other that lies at the heart of the epistemology of silence (whether the relied-upon information source is oneself or another).<sup>7</sup>

Take first the testimony and memory cases. In these cases the relevant source does (purport to) transmit information. In order to have a single expression covering the sort of ‘transmission’ involved in both, it will be helpful to follow Burge (1993) and use ‘presentation-as-true’ and its cognates as the general term covering the positive deliverances of a source. (The idea, then, is that both memory-impressions and attestations/reports are species of the single genus ‘presentation-as-true’.) Here the conditional relevant to assessing the epistemic status of a subject's belief, formed through accepting the presentation-as-true in question, is what we might call the presentation-to-truth conditional:

(1) If source  $\Phi$  were to present-as-true that  $p$ , then it would be the case that  $p$ .

(As stated, (1) is a bit awkward. More intuitively, the claim is this:  $\Phi$  would present-as-true that  $p$  only if it were the case that  $p$ .) If the case is one of memory, then  $\Phi$  is oneself, and the presentation-as-true takes the form of a

249

seeming recollection (or memory impression); whereas if the case is one of testimony, then  $\Phi$  is the (news) source on whom (or on which) one relies for information of this kind, and the presentation-as-true takes the form of a piece of testimony.

The role of (1) in the epistemology of memorial or testimonial belief is best seen as a claim exploited by the epistemologist in determining the reliability of a given presentation-as-true. In saying that (1) is a claim exploited by *the epistemologist*, I mean that it is *not* something that needs to be entertained by the subject herself in the course of memorial or testimonial belief-fixation. It is arguable that the truth of (1), together with a hearer's counterfactual sensitivity to indications that (1) is false, suffice to render a (memorial or) testimonial belief justified. The subject who acquires such a belief need not draw any inferences from premises like (1), nor even so much as have (1) 'in mind' in the process of memorial or testimonial belief-fixation.<sup>8</sup>

Let us move on to the silence cases.<sup>9</sup> These differ both in the subjunctive conditional involved, and in the role that the relevant subjunctive conditional plays in the epistemology of silence. Begin with the conditional involved. Whereas the relevant conditional in cases involving a presentation-as-true is the presentation-to-truth conditional, (1), in cases of silence, the relevant conditional is one that we might call the truth-to-presentation conditional:

(2) If it were the case that  $p$ , then some source would present-as-true that  $p$ .

(2) captures the idea that there is relevant 'coverage' in the domain in question: all relevant truths will be presented as such.<sup>10</sup> This is what lies behind the reasoning one exhibits when, on the basis of not recalling having seen Jones wear a bright red suit, one concludes that he did not regularly wear any such suit; or when, on the basis of not recalling having come across any report of the discovery of WMDs in Iraq, one concludes that there has been no such discovery to date. Of course, it is not (2), so much as one of (2)'s implications, that is used in our reasoning in such cases. The relevant implication is this:

(2\*) If no source has presented-as-true that  $p$ , then  $\sim p$ .

250

In what follows my claim will be that what I am calling the 'epistemology of silence' is best explored by examining the epistemic status of (2\*), and in particular, by exploring the epistemic grounds one may have for endorsing a conditional of this sort.

Not only do silence cases differ from the presentation cases in the conditional that is involved in epistemic assessment; the role that the conditional plays in the epistemology of the resulting belief differs as well. In particular, whereas (1) is a claim used by the epistemologist during assessment (and need not be exploited by the subject who comes to acquire the belief in question), the relevant instance of (2\*) itself would appear to figure in the very process by which a silence-supported belief is formed by the subject—at least if that process is to be a reliable one. To see this, consider the difference between having evidence for  $\langle p \rangle$ , and lacking evidence for  $\langle \sim p \rangle$ .<sup>11</sup> Without introducing further qualifications, it would be false to say that a lack of evidence for  $\langle \sim p \rangle$  is evidence for  $\langle p \rangle$ . So insofar as one's belief in  $\langle p \rangle$  is acquired through one's recognition of a lack of relevant evidence for  $\langle \sim p \rangle$ , the epistemic goodness of one's belief in  $\langle p \rangle$  would appear to depend on how reliable an indication of  $\langle p \rangle$  it is, when one lacks evidence for  $\langle \sim p \rangle$ .

What is more, it would seem that without background beliefs to justify the hypothesis that the lack of evidence for  $\langle \sim p \rangle$  is a reliable indication that  $p$ , a subject who formed her belief in  $\langle p \rangle$  merely on the grounds of recognizing the evidential lack for  $\langle \sim p \rangle$  will have formed an unjustified belief. This is so whether one's theory of justification is reliabilist or more 'internalist' and reasons-based. If one is a reliabilist, then the belief in question is unjustified because unreliably formed; whereas if one is more internalist, then the belief is unjustified since grounded in an inference from ignorance (which is a fallacious inference-type).

Under what conditions, then, might a silence-supported belief enjoy a doxastic justification? To some extent, our answer depends on our favoured theory of justification. If one is a reliabilist, reliability in the formation of such a belief would appear to require having inferred the belief in  $\langle p \rangle$  in a conditionally reliable way from a reliably-formed belief in the relevant instance of (2\*). If one regards justification as a matter of good reasons, one will regard silence-supported beliefs as justified only if the subject has good reasons for regarding the relevant state of ignorance as indicative of the truth of what she believes. But whatever one's favoured theory of justification, we can now see that the acquisition of a silence-supported belief that  $\sim p$  is best seen as a kind of inference from ignorance; and so the role of (2\*) in this process is that of the key premise in the inference from which the subject herself draws the conclusion that  $\sim p$ . Insofar as the subject formed her belief that  $\sim p$  in this way, this belief is justified only if she is justified in believing the relevant instance of (2\*). We have reduced

the epistemology of silence to the question, What does it take to be justified in believing the relevant instance of (2\*)? I propose to return to address this matter below.

3.

Let us first ask after the *truth-likeness* of a belief formed on the basis of the silence of a relied-upon source of information: how likely is it that a belief of this sort will be true? And let us deal first with the case where the relied-upon source is oneself, and the ‘silence’ is that of one’s lacking any relevant memory (CASE 1). As anticipated above (in our discussion of the sources of possible error in these beliefs), two main factors are relevant to the truth-likeness of these beliefs: (a) the *coverage-reliability* of oneself *qua* information source, and (b) the *preservation-reliability* of one’s memory.

Coverage-reliability supervenes on considerations bearing on the subjunctive conditional corresponding to (2) above, as follows<sup>12</sup> :

(S-M) If it were the case that  $p$ , subject  $S$  would have found out that  $p$ .

The idea is that one can rely on oneself for coverage of the relevant facts in the domain in question. Let us say that a  $\langle p \rangle$ -world  $w$  enjoys ‘coverage by subject  $S$ ’ with respect to  $\langle p \rangle$  at time  $t$  when in  $w$   $S$  finds out that  $p$  at some point at or before  $t$ . And call a set of  $\langle p \rangle$ -worlds ‘clean at  $t$  relative to  $S$ ’ when all of the worlds in the set are such that each enjoys coverage by  $S$  at  $t$  regarding  $\langle p \rangle$ , i.e. none of them are such that by  $t$   $S$  does not find out that  $p$ . Then we can characterize a source’s coverage-reliability in a rough but intuitive way in terms of nearby  $\langle p \rangle$ -worlds. The broader the circle of worlds (with the actual world at centre) that are clean at  $t$  relative to the source, the more coverage-reliable that source is regarding  $\langle p \rangle$ . Of course the temporal dimension is also relevant here: if one’s source finds out that  $p$  long after the time by which the subject needed to know whether  $p$ , then such knowledge, even when it is presented to the subject by the source, will be of little or no help. So we might want to restrict  $t$  to the latest time at which possession of the knowledge in question by the subject would satisfy the relevant practical and theoretical aims bound up with her interest in  $\langle p \rangle$ . A source who is coverage-reliable in this sense will be such that for most or all of the propositions  $\pi$  in some domain of interest to the subject, were  $\pi$  true the source would know this, and she would present this knowledge ‘on time’ to help the subject in whatever aims were bound up with the subject’s interests in

252

the matter. Clearly, to the extent that a source is not coverage-reliable, to that extent the source’s failure to discover that  $p$  is a less reliable indication that  $\sim p$ . To just this extent, the beliefs one forms on the basis of the ‘silence’ of that source will not enjoy a high truth-likeness quotient—with the result that it will be increasingly a matter of luck if one’s current ‘silence’-based belief turns out to be true.

Of course, the truth-likeness of one’s ‘silence’-based belief also depends on how well one’s memory *preserves* (the contents of) one’s beliefs. The point can be put in terms of another conditional:

(P-M) If (i) at some earlier point S acquired the belief that  $p$ , (ii) in the interim S acquires no relevant information bearing on whether  $p$ , then if presently S were to be queried regarding whether  $p$ , S would have the memory impression that  $p$ .

A subject's memory may fail to be good in this regard in at least two distinct ways. There is 'straight forgetfulness', in which (although at some earlier point S had acquired the belief that  $p$ ) memory fails to retain that belief. And there are failures of 'content preservation', where (although at some earlier point S had acquired the belief that  $p$ ) what is preserved in connection with that belief is a content distinct from  $\langle p \rangle$ .<sup>13</sup> Either way, as one's memory fails to preserve one's original beliefs (rendering (P-M) false), one increases the risk of error in a 'silence'-based belief. This is easiest to see in cases of straight forgetfulness. Here it should be clear that the more straightforwardly forgetful a subject is, the less her having no recollection of  $\langle p \rangle$  (when queried) correlates with the falsity of  $\langle p \rangle$  (= the truth of  $\langle \sim p \rangle$ ). Arguably, a similar point holds in cases involving the failure of content preservation. Here it can happen that the subject (when queried) has no recollection of it being the case that  $p$ , under conditions in which she originally acquired the belief that  $p$ : her memory shifted the content of her original belief from  $\langle p \rangle$  to (logically distinct)  $\langle q \rangle$ , and her current memory-sustained belief that  $q$  failed to trigger the recollection that  $p$ .<sup>14</sup> Once

253

again, a subject whose memory exhibits this sort of failure is a subject facing an increased risk of falsity in her silence-supported beliefs.

It is clear that these same points hold in the interpersonal case, when the subject relies on the coverage-reliability of some source other than herself (CASE 2). Here coverage-reliability depends on the interpersonal analogue of the subjunctive conditional (S-M), as in the following:

(S-T) If it were the case that  $p$ , one of the sources on whom S relies would have found out that  $p$ .<sup>15</sup>

Suppose our subject forms the silence-supported belief that  $\sim p$ , under conditions in which (S-T) is false. Then her silence-supported belief is exposed to an increased risk of falsity: even assuming that her failure to recall having been told that  $p$  accurately reflects that her sources did not in fact report or publicize that  $p$ , her silence-supported belief that  $\sim p$  is not properly sensitive, since the failure of her sources to report/publicize that  $p$  reflects a failure on *their* part to have discovered that  $p$  (as when the fact that  $p$  has been successfully covered up by the relevant authorities).<sup>16</sup> Alternatively, the truth-likeness of one's silence-supported belief is decreased under conditions involving unreliability in the transmission of the message from source to subject. In these cases the following subjunctive conditional (corresponding to (P-M) in the memory case) would be false:

(P-T) If the source(s) on whom S relies had acquired the information that  $p$ , they would have reported or publicized that  $p$ , in such a way that S would have come across the

transmission (or something that stood in a reliable communication-relation to that transmission, and which preserved its content).

Even on the assumption that one's source(s) is/are coverage-reliable in the sense of (S-T), S's silence-supported beliefs (formed through relying on said source(s)) will have an increase in the likelihood of falsity if said source(s) doesn't/don't reliably transmit the information in question. For in that case, even assuming one's failure to recall having been told that  $p$  accurately reflects that one's sources did not in fact report or publicize that  $p$ , even so, one's silence-supported belief that  $\sim p$  is not properly sensitive, since the failure of one's sources to report/publicize that  $p$  reflects a failure on their part to have reported what they had, in fact, discovered.

In the light of the foregoing considerations about the reliability of silence-supported belief, we can now explicitly address the question regarding the conditions on the doxastic justification of such beliefs. Insofar as the subject

254

formed her belief that  $p$  by inference from the relevant instance of (2\*), her belief is justified only if she is justified in believing that the relevant instance of (2\*) holds. This does not require the truth of that instance of (2\*), so much as it requires whatever is necessary and sufficient for justified belief in (what is expressed by) the relevant instance of (2\*). Assuming a reliabilist theory of justification, this will require that the process or method through which a hearer arrived at her belief in the relevant instance of (2\*) is globally reliable.

Under what conditions would one count as globally reliable in the formation of the belief that

(2\*) If no source has presented-as-true that  $p$ , then  $\sim p$ ?

Although it might be too much to ask for necessary and sufficient conditions on the reliable formation of such a belief, perhaps we can be satisfied here with sufficient conditions. A subject S reliably forms a belief in (2\*) so long as S's own relevant expectations for the news are well calibrated to the sorts of news reports that are prevalent in her community.<sup>17</sup> Other accounts of justification, of course, may demand other things of her in order for her belief in (2\*) to be justified. But whatever one's favoured theory of justification, the point remains: whereas our subject cannot acquire silence-supported *knowledge* unless (2\*) is *true*, doxastic justification in these cases does not require the truth of (2\*) so much as it requires that the subject's belief in (2\*) be justified.<sup>18</sup> Insofar as testimonial belief does not require that the subject have any belief in the presentation-to-truth conditional (1) (let alone that such a belief be justified),<sup>19</sup> testimonial belief and silence-supported belief receive different epistemological treatments in this respect.

4.

Parallels between testimony and memory have been profitably explored elsewhere.<sup>20</sup> It is popular to assimilate the two by treating one on the model of the other, for example, by thinking of memory as a matter of one's 'past' self testifying to one's 'present' self. Though I have touched on this, the plausibility of such an assimilation has not been defended here. My claim, rather, has been that there is another dimension to this assimilation—one not often described by those who liken memory and testimony—and that this dimension is one that brings

255

out in a particularly vivid way another form taken by our epistemic reliance on our social peers. In this connection I have been stressing the similarities between beliefs in  $\langle \sim p \rangle$  formed through noting that one does not remember having *observed* otherwise, and beliefs in  $\langle \sim p \rangle$  formed through noting that one does not remember having *heard* otherwise. With this comparison in mind, we might wonder whether the epistemology of silence-supported belief, in those cases involving our reliance on sources other than ourselves, is fruitfully thought of as the epistemology of 'socially extended memory'.

I begin with the least controversial of the claims that can be made on this score (which is not to say that the claim is uncontroversial!). Epistemology of silence cases involving reliance on sources other than oneself illustrate the following hypothesis in social epistemology: there are occasions on which the epistemic assessment of (the cognitive states of) a given individual will be incomplete without taking into account the dispositions and cognitive states of the individuals' social peers. The illustrating cases are those in which the members of the individual's community play a role akin to that of a perceptual and memory system: they acquire information from the world through perception, they store this information, and they transmit this information in such a way as to enable its subsequent retrieval *by the subject*. It is clear that cases in which the subject's peers play these roles are importantly different, and in many ways, from cases in which only her own perceptual and memory systems are in play. Even so, for at least some of a subject's beliefs, if we are to give an accurate assessment of the epistemic status of these beliefs, then the (perception- and memory-like) roles played by the subject's peers must be taken into account, in a way that is analogous to the way that we take into account the performance of her own perceptual and memory systems, in cases where she relies on these. I call this 'the Thesis.'

If we endorse a broadly reliabilist approach to epistemology, then the Thesis is true.<sup>21</sup>

Suppose that you are not particularly attentive to the way people dress, and that you do not find anything noteworthy in what others would regard as outrageous attire. So, had it been true that Jones regularly wore a bright red suit, even so you would not have attended to (or registered) this fact. Were you then to form the belief that Jones did not regularly wear a bright red suit, on the basis of the fact that you can recall having seen no such thing, your belief would not be reliable. On the contrary, the lack of coverage-reliability



of your perceptual system in the relevant domain translates into a lack of reliability in the belief you form through reliance on that coverage-reliability.

Consider now the case involving epistemic reliance on the coverage provided by one's community. The less coverage-reliable are the individuals on whom one is relying (relative to the domain of facts in question), the less reliable will be

256

any silence-supported belief formed through reliance on the coverage-reliability of these peers. Suppose that neither the *New York Times* nor CNN nor any of the other media sources on whom you rely for the news is particularly attentive to what is going on in Iraq.<sup>22</sup> We can suppose, contrary to the facts, that without publicizing this fact the media have left Iraq after deciding not to put any resources into reporting from there any more, and after deciding not to get any reports from others (such as the Associated Press) who remain there. So, had it been true that WMDs were found in Iraq last week, even so they would not have attended to (or registered) this fact. Were you then to form the belief that WMDs have not been found to date in Iraq, on the basis of the fact that you cannot recall having come across any report of such a finding, your belief would not be reliable. On the contrary, the lack of coverage-reliability of your sources in the relevant domain translates into a lack of reliability in the belief you form through reliance on those sources for coverage.

The point common to both cases would appear patent: the epistemic significance of 'silence'—the epistemic significance of the fact that you recall no source having presented  $\langle \sim p \rangle$  as true—is in direct proportion to the quality of coverage by, and transmission through, the source on whom you are relying. This is as true when that source is someone other than yourself—as in the case of coverage by members of news organizations, where  $\langle \sim p \rangle$  (were it presented at all) would be presented via testimony—as when it is you yourself—that is, in memory cases, where  $\langle \sim p \rangle$  (were it presented at all) would be presented through a memory impression. As we saw above, in both sorts of case the main threat of unreliability in the silence-supported belief can be decomposed into the threats of (a) unreliability in the coverage (= failure to satisfy the truth-to-presentation conditional, (2)), and (b) unreliability in the transmission of the source's information (= failure to satisfy one or both of the 'transmission' conditionals (P-M) and (P-T)). Admittedly, in cases involving reliance on the coverage provided by someone other than oneself, cases are complicated by the fact that, in addition to one's source having to satisfy the transmission conditional (P-T), one must currently be reliable in one's judgement that one did not encounter any such report in the past. But this only suggests that the case of relying on another, in addition to having the same features as cases involving reliance on oneself, has other features besides. This should come as no surprise: self-trust appears to be a basic form of trust.<sup>23</sup>

Might we then make the stronger claim, that silence-supported beliefs formed through reliance on the coverage-reliability of others is akin to a case of 'extended

memory'? In defence of this idea, we might try comparing the role of others as repositories of information with the role of what might be called personal mental prosthetics (diaries and notebooks).<sup>24</sup> The core contention of Clark and Chalmers (1998) is that a diary could count as a case of 'extended memory' so long as (a) it stores information that one has acquired and information about plans for the future that one has put into it, and (b) this information stored there is readily accessible to the subject.<sup>25</sup> Those philosophers convinced by the case from Clark and Chalmers (1998) might try seeing the role other people play in a subject's acquisition of a silence-supported belief as that of a repository of information for the subject, a sort of 'social' version of the 'extended memory' hypothesis.

Unfortunately, the 'extended memory' proposal runs into difficulties with respect to both (a) and (b). It runs into difficulties with respect to (a), in that other people are not repositories for information that *one oneself* has put in those repositories.<sup>26</sup> Of course it might be replied that the assumption, that memory storage is storage of information that *one oneself* has put into the relevant repository, is question-begging against the proposal of a socially extended form of memory. Perhaps a better way of thinking of a subject's memorial storage is as a repository of information that she can exploit in the course of her belief-fixation and belief-revision. So long as information is stored in a repository in such a way that the information is sufficiently accessible to her, and in particular is accessible as needed for belief-formation and belief-revision, one's peers might constitute such a repository (at least by the standards of Clark and Chalmers 1998). Or so it might be argued.

But this brings us to the problems that the 'extended memory' proposal faces regarding (b): it would seem that the information in the minds of those on whom one is relying for coverage is *not* sufficiently accessible to one. In this respect it is worth bearing in mind that there is one important feature of the silence case involving reliance on others for coverage: the subject does not have access to the information stored in others' minds, but instead merely notes that she (the subject) cannot recall having received a relevant report from any of them. The result is that if we regard them as storing information for the subject, this is information that will forever be inaccessible to her (unless she can read their minds or they testify to her). It thus seems that we do not properly regard information stored in others' minds as relevantly accessible to the subject herself.

The considerations just advanced suggest that the case has not been made for regarding our reliance on others for coverage on the model of a 'socially extended memory'. Still, we must find some way to acknowledge the role that others play in the epistemology of silence-supported beliefs. If it is not by way of assimilating their role into that of an 'extended' sort of memory, how should we characterize this role?

Here is a proposal that I favour. (I offer this in the spirit of a suggestion that is worth developing further.) We begin by noting that, in cases where a subject is assuming the coverage of others, reliable silence-supported belief is not the cognitive achievement of any single individual. We then go on to hold that the relevant social facts that underwrite the reliability of a given silence-supported belief—the social facts that ensure that the subject has reliable coverage in the case at hand—are all-important *background conditions* on the formation of silence-supported belief. In particular, it is against these conditions that we assess the epistemic features of silence-supported beliefs. There is in this way a kind of division of epistemic labour. The subject herself must be justified in believing—must have reliably acquired a belief—in the relevant instance of (2\*). This is a matter of her news expectations being properly calibrated to the relevant news-dissemination practices in her community. But whether or not she is properly calibrated in this way, whatever the prevalent news-dissemination practices in her community happen to be, it is against the background of such practices that we assess the reliability of her belief in (2\*), as well as the reliability of her silence-supported belief (inferred from her belief in (2\*)). This sort of view thus regards the cognitive achievement of silence-supported knowledge and justified belief (in cases where the relevant coverage is assumed to be provided by others) to be an achievement that is made possible by the existence of certain social arrangements, even as it denies that we should regard those arrangements as constituting anything like a ‘socially distributed memory system’.

It is worth noting that the proposal to regard social factors in this way is thoroughly in keeping with the spirit of reliabilist epistemology. After all, reliabilists have long assessed the reliability of a belief-forming process against what Hilary Kornblith calls ‘the natural environment in which [the process itself] operate[s]’. Here is Kornblith’s extended comment, made in connection with social factors:

Just as we wish to know whether our natural inferential tendencies are likely to give us an accurate or a distorted picture of the world, we also need to know whether our social institutions and practices are helping to inform us or to misinform us. And just as we need to examine our perceptual and inferential equipment against the background of the natural environment in which they operate, we also need to investigate these mechanisms against the background of the social environment in which they operate. *Such investigations are straightforward extensions of the naturalistic project in epistemology.*

(Kornblith 1994b: 97, italics added)

Since Kornblith makes this point in a paper which is otherwise advocating a ‘*conservative extension* of naturalistic epistemology’ (1994b: 94, italics added), it seems that the present proposal should be relatively uncontroversial to anyone who is already on board with reliabilist epistemology.

No doubt, there is great variation in social practices and institutions worldwide, and even between local communities there can be relevant differences. But this fact of social

variation is no objection to the proposal to regard the relevant social conditions as background conditions on silence-supported belief. For one thing, social variation itself is a normal feature of our larger world-environment. For another, it is a normal part of socialization within a community that one becomes sensitive to prevailing social arrangements. Included in this socialization process is the process of learning what counts around here as newsworthy, what sorts of news one can expect (and where and when one can expect this news), and so forth.<sup>27</sup> Acknowledging all of this merely forces us to acknowledge that the background conditions that obtain in any given case depend on the community in play.

In short: it may be too much to regard silence-supported belief, where one is relying on others for coverage, as involving a kind of ‘socially extended memory’. At the same time, we can recognize the all-important role that a subject's peers play in ensuring the reliability of her silence-supported belief in another way: we can think of the relevant social practices as part of the background conditions on the process(es) used in the formation of silence-supported belief.

5.

The epistemic significance of the silence of a relied-upon source of information is in many respects the same, whether that source is oneself (as in memory cases) or members of one's social community (as in cases of coverage-reliance on a news organization). To say this is to encourage the parallel between memory and testimony, in ways that go beyond the standard parallels. It is also to encourage novel thinking about testimony itself. In addition to relying on what others tell us, we also rely on what they do not: their silence can back our claims to know, and in any case brings to the fore our reliance on them for coverage. I submit that the phenomena of coverage and silence-supported belief go alongside

260

that involving actual testimony, as manifesting distinct varieties of our epistemic reliance on our social peers.<sup>28</sup>

261

## 13 Epistemic Circularity and Epistemic Incommensurability

Michael P. Lynch

### 1. INTRODUCTION

If I were to believe you are trustworthy just on your say-so, my reasoning would be infected with what is called *epistemic circularity*. I would be supposing a source is trustworthy by relying on that very source. Generally speaking, we tend to think this is a very bad idea. It is why we don't bother asking politicians or salesmen whether they are honest.

A well-known line of reasoning stemming from the Ancients—what we might call the criterion argument—claims that epistemic circularity leads to radical scepticism. Baron Reed has recently put the point this way:

Let  $F_1$ ,  $F_2$ ,  $F_3$ , etc., be a subject  $S$ 's cognitive faculties, of which  $S$  has a finite number. In order to know that  $F_1$  is a reliable source of knowledge,  $S$  will have to use either  $F_1$ , or another faculty. But if  $S$  uses  $F_1$  his belief that  $F_1$  is reliable will be epistemically circular. So,  $S$  must instead use (say)  $F_2$ . But  $S$  should not use  $F_2$  unless she knows that it is a reliable source of knowledge itself. In order to come to know this,  $S$  will have to use  $F_1$ ,  $F_2$ , or some other faculty. But  $S$  cannot use  $F_2$ , on pain of epistemic circularity. And  $S$  cannot use  $F_1$ , without first knowing that it is a reliable source of knowledge, which is still in question. So,  $S$  must use some other source—say,  $F_3$ . But it should be clear that the same issues will arise with respect to  $F_3$ , and that  $S$  will eventually run out of faculties to which she has not already appealed. Epistemic circularity is inescapable.

(Reed 2006: 186–7)

The invited conclusion: we don't know whether our faculties or sources of belief are reliable. And if we don't know that, then, according to the sceptic, those same sources cannot be sources of knowledge—that is, they cannot produce justified beliefs. Consequently, says the sceptic, I don't—and perhaps can't—know anything.

262

The criterion argument is disarming in its simplicity. But it is open to an equally simple response grounded in epistemic externalism.<sup>1</sup> The externalist rejects the key assumption of the argument—namely, that one must first know that a source of belief is reliable before it can be said to produce knowledge. From the externalist's viewpoint, all that matters is whether the source *is* reliable, not whether we know or even believe that it is. As long as the source is in fact reliable, it can produce knowledge, including knowledge about whether it, or some other source, is reliable. Epistemic circularity may be pervasive, says the externalist, but it doesn't prevent us from having knowledge or justified belief.

The externalist response to the problem of the criterion continues to be the subject of a wide-ranging controversy.<sup>2</sup> One thought is that far from providing an easy victory of scepticism, the externalists' response only reveals the inadequacy of their view, by demonstrating that knowledge is too easy to come by if externalism is correct.<sup>3</sup> To some, the 'if my sources of belief are reliable, then I can know they are' response simply shows that the externalist is missing the point, or talking past the sceptic, or somehow not addressing the original issue.

Despite these doubts, I'm going to assume for present purposes that some version of the externalist response succeeds in answering the criterion argument. My reasoning, such as it is, is that it is difficult to see how that argument *could* be answered without relying on such a response. Nonetheless, I also think that there is something right about our misgivings about the externalist answer to the criterion argument. What is right about those misgivings is that the externalist response has no traction against a different problem, distinct from scepticism. Like the sceptical argument above, this other problem is rooted in part in the issue of epistemic circularity. But unlike its cousin, it is not a problem about whether we in fact have knowledge or are justified in our opinions. It is about rationally resolving explicit disagreement over the reliability of our most basic methods for forming beliefs. Although I will not try to show it here, this latter problem, what I will call the problem of epistemic incommensurability, is arguably the root worry behind the criterion argument, and, in my view, the reason why Chisholm was right to say that the latter was 'one of the most important and difficult of all problems of philosophy' (1982: 61).

The bulk of this chapter will be concerned with clarifying the problem and its nature. In the final section, I briefly sketch a solution.

## 2. EPISTEMIC INCOMMENSURABILITY

Cain and Abel, let's imagine, are having coffee and arguing about the age of the Earth. Abel asserts with great confidence that the earth is a mere 7,000 years old. Cain, amazed, points out that Abel's claim is not justified by the evidence of the fossil record, the best explanation of which is that the Earth is far older. 'Inference to the best explanation from the fossil and historical record can work sometimes' Abel concedes, 'but the *best* method for knowing about the distant past is to consult the Holy Book; it overrides any other competing evidence'. Cain scoffs and rejects the book as an unreliable source for knowing about the distant past; the *only* reliable method, he insists, is to employ a combination of abduction and induction from the fossil and historical record.

As this toy example indicates, disagreements usually start with a disagreement over the facts. But disagreement over facts often turn into disagreement over whose view of those facts is best supported or justified. And sometimes we move still further up what we might call the epistemic ladder: we begin to disagree over how we *ought* to support our

views of the facts, about the sort of evidence that should be admitted, and whose methods more accurately track the truth. When we do that, we are engaged in a truly *epistemic* disagreement: a disagreement over epistemic principles.

By an epistemic principle, I mean a normative principle to the effect that some source or way of forming beliefs has some valuable epistemic status. The valuable epistemic status I'll be discussing here is reliability. So we might say that an epistemic principle is one that says that some doxastic practice or method is reliable. The reason reliability is valuable is that reliable doxastic methods are those that are likely to produce true beliefs. And having true beliefs, I'll assume for present purposes, is a good; it is what we might call a, or even the, epistemic goal.<sup>4</sup>

Doxastic methods are more or less basic. A less basic method would be a blood test, since the reliability of that test would presumably be justified by employing other methods. Plausible candidates for *basic* methods, on the other hand, include inferential methods like deduction or induction and non-inferential methods, like sense perception. It is hard to see how these sorts of methods could be justified solely by appeal to any other method. Consider just one case: induction. If Hume was right, there isn't much hope in justifying the reliability of inductive inference without relying at some point on induction to do so.

But I will not quibble here about examples. It seems likely that there are basic epistemic methods, and I will assume that there are. Let's say an epistemic principle is fundamental when it is about such a method and derivative when it is not.

264

Belief-forming methods, whether basic or not, are more or less restricted in applicability. Deduction is clearly of wide scope—since deductive inference from true premises, if it is reliable, is reliable for forming beliefs of any sort. Not so with sense perception—sense perception is not a reliable means for forming beliefs about mathematics. It is, broadly speaking, applicable to ascertaining the contingent facts about the external world. Still more restricted are Cain and Able's principles: both assert the reliability of methods with regard to a single domain of facts—facts about the distant past.

Let us say that A disagrees with B over some epistemic principle just when A does not believe EP and B does. Thus one disagrees with someone over an EP in this sense when one either disbelieves the EP or withholds belief in it (e.g. when one doubts that it is true). A *overtly* disagrees with B over some EP just when A explicitly withholds assent from an EP B asserts. The paradigm case of overt epistemic disagreement is where assent is withheld because the relevant EP is simply denied. An overt epistemic disagreement is *mutual* just when both sides to the dispute deny an epistemic principle the other asserts.

Overt mutual epistemic disagreement can be deep or shallow. I'll say a disagreement is *deep* when it meets the following conditions:

1. *Commonality*: The parties to the disagreement share common epistemic goal(s).
2. *Competition*: If the parties affirm distinct principles with regard to a given domain,

1. *Commonality*: The parties to the disagreement share common epistemic goal(s).  
those principles (a) pronounce different methods to be the most reliable in a given domain; and (b) these methods are capable of producing incompatible beliefs about that domain.
3. *Non-arbitration*: There is no further epistemic principle, accepted by both parties, which would settle the disagreement.
4. *Mutual Circularity*: The epistemic principle(s) in question can be justified only by means of an epistemically circular argument.

By these criteria, most epistemic disagreements will *not* be deep. In most cases where my methods of belief formation are questioned, Non-arbitration, and hence Mutual Circularity fail—as might happen when one doctor questions another about the reliability of a given test, and the dispute is settled by appeal to other more basic methods. In other cases, subsequent investigation may reveal that the methods in question do not really compete—perhaps because one is more restricted in application than another. Finally, in some cases, the disputants are simply talking past one another—one is aiming at truth, for example, the other at practical success or some such.

Nonetheless, it seems pretty clear that deep epistemic disagreements can and do occur. Cain and Abel's dispute seems a likely example. The two principles involved, again, are:

265

C: Inference to the best explanation of the historical and fossil record is the most reliable method for knowing about the distant past.

A: Reading the Holy Book is the most reliable method for knowing about the distant past.

We can easily imagine that Cain and Abel share the same goal—to ascertain the truth about the distant past of the planet. Moreover, their principles compete. Both practices cannot be the best for determining what is true of the distant past. And assuming that neither thinks that time travel is an option, there appears to be no shared principle that will settle the matter either.

The case plausibly satisfies Mutual Circularity as well. In Cain's case, it is hard to see how he could avoid appealing to a track record argument to justify the reliability of his method. And it is equally hard to see how he could avoid assuming the reliability of that method in accepting the premises of such an argument. For there seems to be no other means (that Cain would accept) which can serve as a check on whether our abductions about the past are accurate. In Abel's case, we can imagine that the circularity might manifest itself in a very different way. He might, for example, give the following argument for his trust in the Holy Book's pronouncements about the past:

God wrote the Holy Book in the distant past.

Everything God wrote in the Holy Book is true.



Therefore the Holy Book is the most reliable guide to the distant past.

Suppose that we ask Abel why he is committed to the first premise, and he responds that the book says that God wrote it. Since this is a belief about the past, it assumes the reliability of the book as a guide to the truth about the past. Abel's argument is not premise-circular: the conclusion is not literally one of the premises of the argument, but it is epistemically circular—it assumes the reliability of a method in an attempt to show that the method is reliable.

Some will wonder if there is an asymmetry to the cases Cain and Abel can make for their epistemic principles. One might think that Cain, but not Abel, will be able to independently confirm the reliability of his method. In fact, however, both Cain and Abel can offer independently confirming arguments. But any such argument will either be itself epistemically circular or push the disagreement back to the applicability of a different fundamental epistemic principle.

Take Cain first. First, recall that the issue in question is whether abduction from the fossil record is the best method for learning about the *distant past*. Unless time travel is an option, *direct* confirmation of the reliability of this method by observation is not an option. That is, we can't check to see whether the method has produced accurate results by direct observation. This fact is part of what will make any track-record argument for the reliability of the method epistemically circular. Moreover, even if Cain were to appeal to certain non-inferential methods (such as sense perception) to *indirectly* bolster the case for

266

the reliability of abduction from the fossil record, these methods can themselves be challenged. And as we've already noted, a defence of the reliability of sense perception, the most likely candidate in Cain's case, will most probably itself be subject to epistemic circularity.<sup>5</sup> As Alston and many others have argued, it is hard to see how we could show that sense perception is reliable without relying on its reliability. And presumably it will not help Cain to protest that his favoured method is merely a local application of a more general reliable method: abduction. For in the imagined case, Abel does not deny the reliability of that method. What he denies is that abduction from the fossil record is the most reliable method for forming beliefs about the distant past.

In addition, note that Abel too might well try to give some independent argument, in addition to the one I've given above, in defence of his principle, one that does not rely on what the book says. He might appeal to the teaching of various prophets, or other sacred texts. Cain will question the reliability of these sources. So Abel might appeal to still more basic, non-inferential methods of belief formation, such as divine revelation or mystical perception of God's actions and will. But here again the question can (and presumably will) be raised about the reliability of such methods. As in Cain's case, it is difficult to see how Abel will be able to demonstrate their reliability without an epistemically circular argument, for it is difficult to see how one could defend one's

claims to reliably speak for, or perceive, God without appealing either to the book, or to mystic perception again.

These points hark back to Reed's argument given at the outset. The general lesson is that where deep epistemic disagreement occurs—and again, it frequently does not—then the disagreement is ultimately disagreement over fundamental epistemic principles.

I emphasize that the Cain and Abel case is an example. You may not agree that it is a deep epistemic disagreement. If you don't, then choose an example you prefer. If you doubt there are any such examples, then reflect on the possibility of an overt debate over the reliability of induction, or sense perception, or any other basic doxastic method.

This is worth stressing. In order for an epistemic disagreement to qualify as deep, there need not be a clash of principles. It only needs to be the case that one side does not affirm a principle that the other side does affirm, and that Mutual Circularity, Competition, Non-arbitration, and Commonality are all satisfied. So overt scepticism about induction, for example, where the sceptic overtly

267

doubts its reliability, qualifies as a deep epistemic disagreement. Commonality is easily satisfied. Competition is satisfied vacuously since the antecedent will be false. Non-arbitration will be satisfied—presumably because, if Hume is right, there is no more fundamental epistemic principle that can justify the reliability of induction. And Mutual Circularity is satisfied because the principle in play is defensible only by epistemically circular arguments. Consequently, overt scepticism about a fundamental epistemic principle like induction is a default case of deep epistemic disagreement. This is a helpful point to keep in mind if you are doubtful about the prevalence of mutual disagreements like that between Cain and Abel. But it is also helpful because it illustrates that overt sceptical challenges are merely a limit case of a broader phenomenon—epistemic disagreement.

Finally, even if deep epistemic disagreements never occur, it is clear that they could. And that is enough to raise the questions with which we will be concerned. The basic issue is simple enough. Where the debate is over basic methods for finding the truth, common ground is peculiarly elusive. It is not as if the parties to the dispute are sharing evidence and drawing distinct conclusions from it. Rather, they are disagreeing over what should be counted as evidence.<sup>6</sup> As a result, there is a symmetry of argumentation in epistemic disagreements. This is because while we don't share the same evidence in cases of deep epistemic disagreement, our epistemic principles are themselves subject to evidence that has *the same structural character*. Neither side can fully justify their principles without circularity; their principles are epistemically incommensurable.

### 3. THE THEORETICAL PROBLEM OF EPISTEMIC INCOMMENSURABILITY

For all I've said so far, one could be excused for thinking that epistemic incommensurability is anything but a regrettable fact about the human condition. After all, I've already conceded that epistemic circularity doesn't bar us from knowing which

epistemic principles are true. Given externalism, I can give an argument for my epistemic principles. If so, then so long as induction is reliable, I can give an inductive argument for its reliability. And that means, that so long as my epistemic principles are true, I can know they are. So epistemic incommensurability

268

may be unfortunate, but if we've already agreed that scepticism doesn't follow, then what about this poses a distinctly *philosophical* problem?

In fact, it poses two problems. The first looks formidable but, is, I believe, actually not so. The second may seem less deep but is in fact more so. I'll take the easy one first. It is meta-epistemic. I'll lay it out in the simplest way I can and then discuss the premises.

1. Deep epistemic disagreements are rationally irresolvable (epistemic incommensurability).
2. The best explanation for why deep epistemic disagreements are rationally irresolvable is that there are no objectively true fundamental epistemic principles.
3. If there are no objectively true fundamental epistemic principles, there are no objectively true derivative epistemic principles.
4. All epistemic principles are either fundamental or derivative.
5. Therefore, there are (probably) no objectively true epistemic principles.

The argument parallels a familiar argument for moral anti-realism. That argument moves from the premise that deep moral disagreements are rationally irresolvable to the conclusion that there are no objectively true moral principles. Such arguments are widely used to undermine moral realism.<sup>7</sup>

As I see it, the first premise of the argument is reasonable, but only when qualified in a certain respect. But once qualified in that respect, the second premise is false.

Let's start with the first premise. What is meant by 'rationally irresolvable' and why believe that epistemic disagreements are so? One relevant sense of 'rational' here is presumably *epistemic rationality*. Epistemic rationality trades in epistemic reasons. An epistemic reason is a reason for thinking that some belief or principle is true. Let us say that a debate is *epistemically resolvable* just when it is possible that an epistemic reason could be *given*, as I'll call it, for resolving it. A debate is epistemically irresolvable when this is not the case.

What do I mean by 'giving a reason'? Suppose we are wondering which of two plans for building a bridge is the safest, and we consult an engineer to find this out. In so doing, we not only expect the engineer to be able to discriminate whether the one plan is safer than the other, but to be able to articulate, or otherwise make clear, the rational basis of her discrimination. In short, we expect her to *give us a reason* for favouring one plan over the other—to articulate the rational basis of her discrimination so that *we can recognize it as such*. For only if we can, from our evidential standpoint as it were, recognize her reason

for favouring one plan *as a reason* can it serve to rationally underwrite our subsequent plan of action. Barring a recognized reason that one or the other plan is safest,

269

we will either not build the bridge at all, or we will judge them equally safe, and choose between the plans in some other way.

This suggests a general constraint: where A gives a reason (in the sense intended) of some type to B for some  $p$ , it must be possible for B to recognize, from his standpoint, that it is a reason. This constraint is consistent with some familiar distinctions. One might be justified in believing  $p$  without having a reason for the belief. The belief may simply be produced by a reliable method. Likewise one's belief might be justified by a reason without you recognizing that it is a reason. But that is not what is at issue in the case of the engineer, or in the case of an overt epistemic disagreement. In both cases, there is a standing demand to *justify* one's beliefs or commitments. And to actively justify (epistemically) a belief essentially requires *giving an epistemic reason* for the belief, and one gives a reason only if it can be recognized as such by oneself and others.<sup>8</sup>

So a debate is epistemically irresolvable just when no epistemic reason can be given for resolving it. Overt deep epistemic disagreements are marked by epistemic circularity. Together, these facts support the first premise, at least when it is understood in the epistemic way we've been developing. This is because epistemically circular arguments cannot be used to give a reason for why some epistemic principle is true.

This is not to deny that an epistemically circular argument might *be an epistemic reason* for believing a principle (the belief might be justified by such an argument). But it cannot be used to effectively give an epistemic reason for a belief in the face of a standing demand to actively justify the belief. For it will fail to be recognized as a reason by anyone who doesn't already share the principle. As the bridge-plan case illustrates, in a case where live demands for justification are on the table, we demand reasons we can recognize as such. Epistemically circular arguments are hopeless for doing this. If I am wondering whether I can trust you, the fact that you say I can doesn't give me a reason for why I should—even if, in fact, I can.

For this reason, epistemic externalism, even if it true, is of little use here. We may well know (via an epistemically circular argument perhaps) which basic methods are reliable. But that fact has absolutely no traction when one is trying to justify employment of a method in the face of disagreement. When posing a challenge, it doesn't help to be told that 'if you adopt the right method, you'll be able to know that you did'. That answer offers stone instead of bread.<sup>9</sup>

270

Nor, I think, is it any more help to claim that one is *entitled*—in the sense intended by Crispin Wright (2004)—to one's fundamental epistemic principles. To be entitled in Wright's sense is, roughly, to have epistemic warrant for a proposition without having any epistemic reasons to believe it. One is entitled to a proposition in this sense when (a)

that proposition is presupposed by one's cognitive projects; (b) one has no reason to think it is not true; and (c) any attempt to justify it would rest on propositions whose epistemic standing is no better off. Perhaps we are entitled in this sense to our fundamental principles. And perhaps, as Wright suggests, entitlement is epistemic and not pragmatic in character. The problem is that even if these points are true, citing that one is entitled to believe some proposition is not going to help rationally resolve a deep epistemic disagreement.

Suppose, for example, that Abel claims to Cain that he is entitled to his fundamental epistemic principles. This may well be so. We can imagine that treating the Holy Book as the most reliable guide to the past *is* presupposed by one of his basic cognitive projects—like understanding the will of God. And he sees no reason to think it is unreliable. Cain of course will hardly be moved, nor should he be. He doesn't accept Abel's cognitive project (he may even think it is incoherent). And from his standpoint he has offered a reason to believe that Abel's methods are unreliable, and therefore that his fundamental epistemic principles are false. But of course, whether this is so is the very issue they are disagreeing about. If their epistemic disagreement is truly deep, Cain and Abel aren't going to recognize each other's epistemic reasons in the domain of belief in question precisely because those reasons themselves are produced by methods that are the very methods in question. In short, 'I'm entitled to *p*' cannot be used to give a reason for *p* in a case of deep epistemic disagreement, *even if* the person who is entitled is warranted in holding *p*.

So understood, the argument's incommensurability premise seems reasonable. Deep epistemic disagreements are rationally irresolvable in the epistemic sense: the principles involved require defence by an epistemically circular argument, but epistemically circular arguments can't be used to give reasons. But of course, it also leaves room for the claim that epistemic disagreements might be rationally resolved in another sense of rationality. I'll return to just that suggestion shortly. For now, I'll bracket premise one and turn to the second premise.

#### 4. A BETTER EXPLANATION

When qualified as I've done above, the first premise of the above argument is reasonable. Deep epistemic disagreements are rationally irresolvable in the epistemic sense. But once qualified in this way, the second premise appears to be false.

271

The second premise is a straightforward appeal to best explanation. In order for it to be plausible, the advocate of the theoretical argument has to rule out the other alternatives – or at least the most reasonable ones. This I don't think can be done. Indeed, I think there is a much simpler explanation available. But first let's briefly consider a few other alternatives.

A *sceptic* will claim that the best explanation for why deep epistemic disagreements are epistemically irresolvable is not that there are no objectively true epistemic principles, but that we can't know which epistemic principles are true. But if externalism about

knowledge is correct, as I've been assuming, we *can* know which epistemic principles are true. So given externalism, scepticism is not the best explanation of the rational irresolvability of epistemic disagreements.

This consideration also underlines an important fact already noted but worth noting again: namely *that the argument is consistent with externalism being true*. For an externalist will claim that if our epistemic principles are true, then we can know they are true (by way of those same epistemic principles). But this conditional can be true even if the antecedent is false. So externalism is not ruled out by the argument.

Of course, if externalism is false, then it may well be that the best explanation for why deep epistemic disagreements are epistemically irresolvable is that we can't know which epistemic principles are true. But that is cold comfort surely.

Another possibility is that premise two is false because the epistemic irresolvability of deep epistemic disagreements is due to defective concepts. That is, we might think that our basic epistemic concepts are simply too vague or imprecise, or otherwise mangled, and that is why disagreements which essentially employ such concepts can not be epistemically resolved. This too would be cold comfort. For if our basic epistemic concepts are radically defective, it may be that many of our epistemic principles which employ such concepts will lack a determinate truth value. We end up reaching the same conclusion by other means.

And of course, there are the familiar possibilities of non-factualism and relativism about epistemic principles.<sup>10</sup> But far from being rival explanations, non-factualism and relativism presuppose the general negative claim (that there are no objectively true epistemic principles) allegedly supported by the above argument. Consequently, it remains our first task to see whether there is a simpler explanation of the epistemic irresolvability of epistemic disagreements.

There is a simpler explanation. The simpler explanation for the epistemic irresolvability of deep overt epistemic disagreements is that it is a direct consequence of the nature of epistemic rationality itself together with the relevant facts about the kind of disagreement in question. More precisely: once one understands what it means to give an epistemic reason, the epistemic irresolvability of *deep* epistemic disagreements is just what one should expect. It is a consequence of the

272

human cognitive condition; what Pritchard has memorably called our 'epistemic angst'.<sup>11</sup>

Where there is deep epistemic disagreement over some fundamental principle, the disagreement has hit bedrock, the spade has turned. It has hit bedrock precisely because the disagreement obeys Mutual Circularity, Non-arbitration, and the other constraints constitutive of deep epistemic disagreement. And when bedrock is reached in this way, the process of giving epistemic reasons—of engaging in the *activity of justifying in the epistemic sense*—loses its point. But the best explanation for this is not some deep metaphysical fact about the objectivity of our epistemic principle. The best explanation is

that epistemically circular arguments for the reliability of some *method won't be recognized as a reason to accept that method by those challenging its reliability in the first place*. In other words, the very nature of the debates themselves, together with facts about what is to give a reason, already explain the epistemic irresolvability of deep epistemic disagreements.

So it is plausible that epistemic disagreements are not epistemically resolvable. But that fact should not incline us to think that there are no objectively true epistemic principles. Nor does it entail that we can't know which epistemic principles are true. It merely demonstrates that, like all things, the process of giving epistemic reasons has its limits, limits imposed by epistemic circularity and the nature of epistemic disagreement.

## 5. THE PRACTICAL PROBLEM OF EPISTEMIC INCOMMENSURABILITY

So deep epistemic disagreements are epistemically irresolvable; but that does not warrant an irrealist conclusion. It may seem tempting to some to leave it at that. But that would be a mistake, for to do so would be to avoid the obvious—and real—remaining problem. For we've just conceded that deep epistemic disagreements are not rationally resolvable in the epistemic sense. That leaves it open that they might be rationally resolvable in some other sense. But what other sense, and how, in that sense, are such disagreements to be rationally resolved?

This seems, to me anyway, a live and pressing question. It is quite different from the question of whether we can answer the Cartesian sceptic. There are very few, if any, radical sceptics. And there are unlikely to be any. But disagreements over basic epistemic methods are, in my view, quite possible. Moreover, the issue at hand is not whether one side or other of an epistemic disagreement can *persuade* the other that their principles are the true ones. Persuasion is a

273

psychological phenomenon, and can be accomplished in a variety of ways. (A big club is a good method; so is an expensive advertising campaign.) What we presumably want here is not a psychological but normative explanation. We want to make sense of what ought to be done in the face of deep epistemic disagreement.

The fact (if it is a fact) that deep epistemic disagreements can't be epistemically resolved means we have to look elsewhere than epistemic reason. An obvious alternative is practical reason. Indeed, one might think that this would have been a reasonable starting place anyway since there is a clear sense in which the problem at issue is practical. Considerations given above, after all, already suggest that the underlying issue isn't a matter of what we know or don't know, but of what we should or shouldn't do. It is not even best conceived as a problem about which principles we 'ought' to believe (because one might say we ought to believe what we know to be true). After all, the root issue at the heart of an epistemic disagreement—that which makes the dispute an 'epistemic' one—is the question of which methods we ought to *employ*. What we want is a reason for employing one method over another. That's a practical matter. No wonder epistemic

disagreements seem irresolvable when conceived of as epistemic matters. Epistemic reasons are simply not the right tool for solving them. It is not a matter of trying to justify our belief in our epistemic principles, it is a matter of trying to justify our actions—our employment of a method.

So I suggest that we reframe the problem of epistemic incommensurability as a practical problem. In saying this, I don't mean that it is a moral problem. Not all questions about what to do are moral questions. The question we are faced with in the face of an overt deep epistemic disagreement concerns our ability to justify our doxastic methods. Since these methods are aimed at producing (true) beliefs, the issue at hand is hardly divorced from epistemic matters. It is practical, but not morally practical. Call it a matter of *epistemic practicality*.

## 6. EPISTEMIC PRACTICALITY

The main goal of this essay has been to get clear on the problem caused by the phenomenon of epistemic incommensurability. That problem, I've argued, is an epistemically practical problem. It concerns how we can justify our employment of our basic epistemic methods in the face of overt disagreement. I've argued that this problem is distinct from the issue of whether we can know that our methods are reliable, and moreover, distinct from the question of whether it is true that they are. Moreover, it is a not a problem that will be resolved by appeal to epistemic reasons. It will be solved, if at all, by appeal to practical reasons.

I'll conclude with a suggestion for how to arrive at such reasons. I say the suggestion is about 'how to arrive at' such reasons because one cannot foresee in advance what specific practical reasons might be relevant to cite in favour of

274

a method for any disagreement over that method. What is needed is a general framework—a way of identifying the types of practical reasons relevant for settling on one method or another.

In setting up the framework, we must recall the general constraints about rational resolution of disagreements uncovered above. Rational resolutions require the giving of reasons. And reasons are given only when they are recognizable as reasons. Together, these constraints suggest that no practical reason for resolving an epistemic disagreement will be successful if it doesn't employ a type of reason—or if, you like, a value—shared by both sides of any such debate. Briefly put, we need common ground.

Of course it is an open question whether such ground can be reached. But rational self-interest provides an obvious initial suggestion. How might we employ such a consideration to argue for one method over another? One possibility is to play what I'll call the *epistemic method game*.<sup>12</sup>

Let us say that an epistemic principle is privileged when it is worthy of teaching in the schools, used in evaluating research, and seen as trumping other, possibly conflicting



methods. In the epistemic method game, players are charged with cooperatively coming up with reasons for privileging some epistemic methods for forming beliefs over others in some world  $w$ —a world distinct from their own. Candidate methods include methods like deduction, induction, sense perception, reading palms, and consulting sacred texts. Players operate, in good Rawlsian style, under the following rules or constraints. First, they cannot, in their deliberations, presuppose that any method is more reliable for producing true beliefs than others in  $w$ . That is, they must operate under the assumption that consulting the sacred texts may well be as reliable a method of belief formation for the inhabitants of  $w$  as sense perception. In effect, their deliberations must take place under the assumption that scepticism is true—all methods for forming beliefs are as reliable as others because none are reliable. Second, just as they cannot, in their deliberations, assume that one method of belief formation is more reliable than any other, so they cannot assume that one metaphysical picture of the world is any more accurate than others. By a ‘metaphysical picture’ I mean—admittedly somewhat loosely—a view about the ultimate structure and nature of reality. Thus, players cannot, in their deliberations assume that, for example, naturalism is true, nor that, for example, Christian theism is true. Third, players know that they will each eventually inhabit  $w$  *themselves*. But finally, fourth, they don't know all the methods they will—because of upbringing, education, religion, and so forth—wish to employ themselves in  $w$ .

Were we to play the method game, it would seem in our self-interest to favour privileging those methods that, to the greatest degree possible, were *repeatable, adaptable, public, and widespread*. Repeatable methods are those that in like cases

275

produce like results. It would be in our interest to favour repeatable methods because such methods could be used over and over again by people with different social standings. Adaptable methods are those that can be employed on distinct kinds of problems and which produce results given a variety of kinds of inputs. It would be in our interest to favour such methods because we don't know what sort of problems we'll face in  $w$ . Public methods are those whose effectiveness could, in principle, be judged publicly—that is, it is not the case that only one person is its sole judge of effectiveness. It would be in our self-interest to favour public methods because we don't know if we'll be lucky enough to be that one person in  $w$ . And finally, widespread methods are those that many people can in fact employ. It seems rational that we would privilege methods with these features simply because by doing so, we would maximize each of our chances to both use and assess the use of the privileged methods. In this sense, such methods could be called democratic.

But aren't players of the epistemic method game using various methods (such as a priori and causal reasoning) to arrive at their conclusion about  $w$ ? Obviously they are, and they may be using other, more idiosyncratic methods as well. The rules do not in any case prohibit this. The situation described is not one where the participants are without epistemic methods, nor is it one where the players must share epistemic methods. It is

one where they are asked to decide—using whatever methods they have available, and acting under the relevant constraints—which methods should be politically *privileged in w*. And the methods that are so privileged are those that will form the content of their subsequent epistemic standards and principles.

One might object that Abel was asking Cain for some reason to believe that certain standards of reason are true. Our suggested form of argument gives something else. It gives practical, self-interested, reasons for adopting some standards over others. Why isn't this stone instead of bread?

Two answers suggest themselves. The first is that stone soup is sometimes the best item on the menu. Given the dialectical situation forced upon us by the nature of deep epistemic disagreements, giving practical reasons is the best we can do. Moreover, practical reasons—and this has been my point—are still reasons. They are better than big sticks.

There is something to be said for the first answer. But there is a better answer available. Namely, not only is giving practical reasons the best we can do, *it is what we should do*. For if Cain and Abel are to engage in real discussion, they should respect each other as fellow judges—as being capable of arriving at the truth. And where we respect each other as fellow judges, we *should*, where possible, defend our claims with reasons that we each recognize as reasons. And that means, if we are to satisfy the demand to respect each other as fellow judges, we *should* appeal to the common currency of our self-interest to justify why we should privilege certain methods. For only by doing so can we give reasons that can be recognized as reasons. Put most provocatively: if we wish to respect each

276

other as fellow judges of the truth in cases of deep epistemic disagreement, we should not appeal to reasons for judging something to be true.

I have not tried to argue *which* specific methods would emerge from the epistemic original position, although I very much doubt that ‘reading the sacred text’ would be one. My point is that *even in the face of disagreement*, we can give reasons for our methods. The reasons are practical, not theoretical, but they are reasons all the same. Moreover, we should try to give such reasons. If we are to treat each other as autonomous judges worthy of equal respect, we must engage in the process of giving and asking for reasons, even when the question at hand concerns the reliability of our most basic methods for reaching the truth.<sup>13</sup>

277

## 14 The Epistemology of Disagreement

Ernest Sosa

This chapter will take up one main question: *when and how can a belief be sustained reasonably in the face of known disagreement?* On the answer to this question will depend our prospects for sustaining cherished opinions in fields where controversy abounds, such as religion, politics, history, morality, art, philosophy, and even medicine and the law.

1.

We begin with a subsidiary question: is reasonable disagreement ever possible? Opposing answers to one and the same question can both be reasonable, of course, if at least one of them is based on evidence that is persuasive but misleading. This much is uncontroversial.

In a more interesting case, Pro and Con share all their evidence. Can they still assess the shared evidence differently? Can one affirm what the other denies, though each proceeds reasonably enough? For each to be reasonable, each needs positive justification. Unlike ethics, epistemology repels arbitrariness. Facing a choice between bringing it about that *p* and bringing it about that not-*p*, you may have no sufficient reason to prefer either over the other, in which case you might well be free to take your pick. That's how it is for practical choices or actions. By contrast, with no more reason for *believing* either a proposition or its negation in preference to the other, you are definitely not free to proceed either way. Here you *must* withhold, if you are to proceed reasonably at all, *epistemically*. If two opponents are both to be reasonable, then, each needs a balance of reason favouring his side.<sup>1</sup> But is this compatible with their sharing all of their evidence?

278

Not if any reason they may have, for or against believing, would have to be found in the evidence that they share. We are supposing they share *all* their evidence. Since the evidence cannot point in two opposite directions at once, Pro and Con cannot each have substantial positive reason for affirming what the other denies.

Based on such reasoning, you may well conclude that reasonable disagreement with full disclosure is just impossible. But others will no doubt disagree. Even after you pool your evidence, suppose they still remain unimpressed.

On one view, with substantial support in the literature, if you encounter opposition from an apparent peer, then, absent *independent* reason to downgrade him, you must lower your confidence, perhaps below the threshold of belief. 'Without independent reason to downgrade an overt opponent', we are told, 'the only reasonable option is to withhold'. Upon encountering opposition this view is self-defeating. Any subscriber to it who is unable to demote independently an opponent who disagrees is then required to withhold.

His opponent, by contrast, labours under no such handicap, and is able coherently to sustain his side.

Enemies of reasonable disagreement advocate an even stronger thesis, moreover, the Uniqueness Thesis:

No body of evidence E justifies more than one doxastic attitude (believing, disbelieving, withholding) on any given question.<sup>2</sup>

When Pro and Con fully share their  $\langle p \rangle$ -relevant evidence E, we are told, each must withhold unless he can find some independent reason for demoting the other.<sup>3</sup>

279

Consider the following assumptions:

1. A pondered question, whether  $p$ , admits just three doxastic attitudes: belief, withholding, and disbelief. Which of these one should adopt is determined by the total  $\langle p \rangle$ -relevant evidence in one's possession, along with the fact that it is one's total evidence.
2. 'Full disclosure' between Pro and Con, on the question whether  $p$ , results in their sharing some E as their  $\langle p \rangle$ -relevant evidence, in such a way that E is for each the total evidence bearing on that question.
3. One and the same total evidence E cannot favour both  $\langle p \rangle$  and  $\langle \text{not-}p \rangle$  at once. The degree to which it favours  $\langle p \rangle$  is the same as the degree to which it disfavors  $\langle \text{not-}p \rangle$  and vice versa.
4. Whether S is reasonable to believe that  $p$  depends on the degree to which his total evidence E favours  $\langle p \rangle$ . E must favour  $\langle p \rangle$ , and must do so sufficiently.

It follows from these assumptions that Pro and Con cannot both be reasonable, with full disclosure on a certain question, when one believes what the other disbelieves. It is not possible to disagree reasonably under full disclosure.

Reasonable disagreement can be rife even so, especially in domains like politics, religion, and philosophy. For one thing, such disagreements rarely take place under *full* disclosure. For another, it is not obvious that what matters here is always (ever?) just a matter of *evidence*, properly so called, much less of *disclosable* evidence.

Perhaps the main reason why disagreements on controversial issues are often reasonable on all sides is that they are to some extent verbal. Verbal but not uninterestingly superficial, as when someone believes there's a 'bank' nearby while someone else 'contradicts' him. Disagreement can be reasonable *and interesting* when its nature as verbal can emerge only through patient, extended analysis and discussion.

Such disagreement can be more or less verbal. Here's one way to think of that. Suppose analysis reveals that sentence X, as understood by S, is true if and only

if  $p$ , while as understood by  $S'$  it is true if and only if  $p'$ . Then we can wonder about the extent to which  $\langle p \rangle$  and  $\langle p' \rangle$  overlap in content. And this might of course vary enormously.

At one extreme the overlap is minimal, as when for  $S$  the meaning is that a river bank is near, whereas for  $S'$  it is that a financial institution is near. But the overlap can be huge, as where for  $S$  the sentence  $X$  is tantamount to an  $n+1$ -termed conjunction  $\langle p_1 \dots \&pn\&q \rangle$ , whereas for  $S'$  the sentence is tantamount rather to  $\langle p_1 \dots \&pn\&q' \rangle$ . If  $S$  affirms  $X$  while  $S'$  denies it, but each would agree on the conjunction  $\langle p_1 \dots \&pn \rangle$  then their verbal disagreement may be as largely verbal as is the disagreement of those who disagree on the sentence 'There's a bank nearby'.

Still, though largely verbal, such disagreement need not be superficial, as is the one about banks. How *substantive* (and not merely verbal) a disagreement is in this sort of example will depend on how many of the  $n$   $p$ -conjuncts the parties disagree about, with some weighting presumably for the 'substantive importance' of the various conjuncts. How *deep* (non-superficial) their disagreement is will depend on how much work of conceptual analysis would be required in order to uncover its real locus, and thereby whether it is really as verbal as is the one about banks.

Though convinced that much disagreement on controversial issues, especially in philosophy, has that deeply verbal nature, I here put that aside, because (a) it deserves a discussion of its own (preferably with philosophical case studies, such as that of disagreement among epistemologists on issues concerning epistemic 'justification'), and because (b) I discuss it elsewhere, in defending armchair philosophy from attacks on the evidential worth of intuitions when people ostensibly disagree in their ostensible intuitions.<sup>4</sup> Here we shall pursue other ways in which ostensible disagreement might be compatible with beliefs that remain reasonable even when they ostensibly clash. How might one reasonably defend one's own side in such a reasonable disagreement?

2.

Several philosophers have converged on the view that the proper response to disagreement among ostensible epistemic peers is for each to give equal weight to the opponent's view.<sup>5</sup> We are told *that it would be unreasonable to downgrade the opponent based simply on the substance of the disagreement*. Here are some examples offered in support of that claim.

1. Pro looks out on the quad and thinks ‘There goes the dean walking on the lawn again’. The day is bright, the quad treeless, the lawn visible throughout, from edge to edge. Yet Con, standing next to Pro, intones ‘It’s nice to see the quad with no one there at all’. What is Pro to do, once it is clear that Con is dead serious? One thing he cannot properly do, according to our converging philosophers, is to downgrade Con’s epistemic credentials based just on their disagreement.
2. At the end of a restaurant meal, Pro and Con mentally compute and Pro arrives at \$42 as what they each owe, whereas Con comes up with \$45 as the right amount. One thing that Pro cannot properly do, in sustaining his side, is to downgrade Con’s epistemic credentials on the question before them, based just on their disagreement.

To demote your opponent based just on your disagreement is likened to declaring someone else’s watch inaccurate because it disagrees with your own, in the absence of any independent basis to prefer either watch over the other.<sup>6</sup>

If I can find no independent reason to downgrade an opponent below my epistemic level, I must consider him no less than a peer. What does it take to be someone’s peer on a given question? Some require parity on general competences that bear on how to answer that question, plus familiarity with the pertinent evidence and arguments. Alternatively, one might define A and B as peers on a given question if and only if, conditional on their disagreement, they are equally likely to be right on that very question.<sup>7</sup>

282

Our two definitions are intimately related *if* how likely you are to be right on a given question is determined just by your relevant competences and your familiarity with the pertinent evidence and arguments. However, the likelihood of your being right on the question may be determined by more than just *having* such competences and familiarity. It may also depend, for one thing, on how likely you are to *employ* them. In order to decide whether these factors are required *in addition*, we need a fuller conception of competences and of how they relate to the relevant arguments and evidence. When facing a certain question, for example, would you necessarily be lacking (somewhat) in competence if you failed to employ your competences on the relevant evidence and arguments familiar to you?

Pending such refinements, the account before us seems reasonable enough as a working hypothesis. In any case, the interesting question here concerns what it would be reasonable to adopt as one’s epistemic stance when one disagrees with someone respected as a peer.<sup>8</sup>

Suppose you share with your friend all the available evidence concerning a shared question. Suppose at first you count your friend a peer, and then it turns out you disagree. According to the Equal Weight View, *when you learn of your friend’s disagreement, you are called upon to consider him equally likely to be right.*<sup>9</sup>

We shall discuss that view in due course. But first we consider its implications for avoiding a dispiriting ‘spinelessness’ whenever we disagree with apparent peers. The

importance of this issue is shown by how broadly it would bear on questions that matter to us in politics, religion, philosophy, and other such

283

domains. It would be bad to have to suspend judgement on just about any controversial question.<sup>10</sup>

Some have urged that we must be adults and face the facts. This we may be forced to do, but a more hopeful option first demands attention: why not protect our controversial views by denying that we are likely to consider someone a peer on an important and controversial issue if we disagree with him on that issue.<sup>11</sup>

Let the issue be abortion, for example, on which Ann and Beth deeply and persistently disagree. Ann surely does not consider Beth a peer on that issue. This is because she takes Beth to be profoundly mistaken on many linked issues. Once Ann sets aside Beth's opinions on a host of issues linked with that of abortion, Ann may simply not consider Beth a peer.<sup>12</sup>

On an issue that divides them, Pro might take Con to be either his peer, or inferior, or superior. On the other hand, Pro might consciously suspend judgement on any or all of these, and in particular might consciously suspend on whether he, Pro, is superior to Con on the question at hand.

With these distinctions in mind, we return to Ann and Beth on abortion. Here is the proposal before us:

It fails to be the case that Ann will consider Beth a peer on whether abortion is permissible, once so many of Ann's considerations have been set aside.

Nor does Ann consider herself superior to Beth, presumably, for the same reasons. So, here's Ann's state of mind: She is sure that abortion is permissible in a given set of circumstances. But she does not consider herself *more* likely to be right than Beth, who openly *denies* that abortion is there permissible. Ann fails to consider herself superior to Beth on that question. But she now seems incoherent. Ours is hence after all no way for Ann to sustain her positive opinions in the face of controversy. She needs not only to avoid taking her opponents to be peers; she needs somehow, more positively, to consider herself superior.

Not only must one *fail* to consider an opponent at least equally likely to be right. That is perhaps Ann's position with respect to Beth. As has been suggested, Ann may be unable to pass judgement on Beth's credentials concerning the question of abortion, given the gulf that divides them, since their respective linked opinions differ so radically. Let us grant this for the sake of argument.

284

True enough, one will thereby avoid the incoherence of considering oneself reasonable in taking abortion to be permissible despite granting that an epistemic peer believes it to be

impermissible. Nevertheless, an incoherence still looms. While facing the question of how likely one's opponent is to be right, one is now forced to withhold on that question; at least one is forced consciously and deliberately *not* to render a judgement. This puts one in the following position: (a) one thinks that  $p$ , and that one is likely enough to be right in so thinking, and (b) one also takes someone else to think that  $\text{not-}p$ , while (c) one consciously fails to have a view on how likely this opponent is to be right. Doesn't simple rational coherence require us to be more judgemental than that? If we think ourselves likely to be right in believing that  $p$ , and we know our opponent to disbelieve what we believe, are we not constrained to consider that opponent likely to be wrong, as likely to be wrong as we are to be right? We fall short, surely, if we consciously remain neutral on whether the opponent is right. Such neutrality fits ill with insisting on our own positive view.

The fact that deep and broad disagreement will preclude one's respecting an opponent as a peer fails therefore to protect us from the dispiriting implications of the Equal Weight View. Avoiding those implications requires rather downgrading one's opponents as somehow inferiors (on the questions in dispute).<sup>13</sup> If we wish to avoid those implications, therefore, we need to reconsider the Equal Weight View itself.

Here now is a capsule statement:

EWV Absent *independent* reason to downgrade an opponent epistemically on the question whether  $p$ , one must give their opinion equal weight as one's own. (One cannot properly insist on one's side of the disagreement while drawing from the very fact that someone disagrees *on that question* a reason for considering their relevant judgement inferior to one's own.)

This applies with special plausibility, so it is contended, in any case where disagreeing parties have fully disclosed their evidence, and where neither has independent reason to downgrade the other's ability to assess the now fully shared evidence.

With such disagreement out in the open, we turn reflective, and consider what our evidence might be, and what it might render plausible or better on the

285

disputed thesis. Pro may then find himself convinced that his evidence supports the affirmative on that thesis, whereas Con takes their shared evidence to support the negative. When this happens, then, unless Pro has independent reason for preferring his own judgement, he now has evidence, including Con's opposition, that requires him to suspend judgement.

Why should that be so? We are told that confrontation with an opponent may reveal that we have no *reason* for preferring the view that our evidence supports  $\langle p \rangle$  above our opponent's view that it supports  $\langle \text{not-}p \rangle$ . Nor, often enough, has our opponent, for his part, any *reason* that advantages his side of *this* disagreement. When that happens, as it often does on controversial topics, then the right attitude is to suspend belief. Both sides



of the disagreement should come together in that middle position. Once we must withhold on whether our evidence supports our belief that  $p$ , it is not possible coherently and reasonably to sustain the belief that  $p$ .<sup>14</sup>

3.

Can we really sustain our side of a disagreement only by finding some *independent* supporting reasons, while the opponent has no balancing reasons in favour of his side?

Much will depend on how we conceive of *reasons*. Suppose you have a headache. What reason have you for thinking that you do? The important reason is, quite plausibly, simply *that* you do! Is this a reason that enables you reasonably to sustain your side of a disagreement when an employer believes you to be a malingering faker, with no headache at all. If so, then you can after all demote an opponent by relying on the substance of your disagreement. A huge part of your reason for rejecting the employer's claim that you're faking it is the very fact that gives content to your belief, the fact of the headache itself. Here then one has a conclusive reason that makes one's belief a certainty, even if that reason will be useless in a public dispute. It will not much advance your cause to just assert against your employer that you *do* have a headache, even if this is in fact the reason that makes you certain that you do. (This gulf between private and public domains will come to the fore again in due course.)

Nor are headaches special in this regard. The same point applies to any obstreperous enough mental state. Anyone who denies to you that you are in that mental state is in a position like that of the employer who accuses you of faking it.

286

Other examples of the same epistemic phenomenon will be found in any case of the given, whether it is the phenomenal given, such as our headache, or the rational given, such as the simplest truths of arithmetic, geometry, or logic. Here again if someone denies what you affirm, you can uphold your side by appealing to the very fact affirmed. Thus, if someone claims 2 and 2 not to equal 4, or a triangle not to have three sides, we could reasonably insist on what we know. We could try to disabuse him, while downgrading his conflicting judgement. *Something* must be misleading him, even if we can't see what that is. And the basis for our stance might well be, as with the headache, just the substance and fact of the disagreement.<sup>15</sup>

That seems to me conclusive on the general issue of whether you can demote an opponent based essentially on the substance of your disagreement, *even when you have no independent reason for doing so*. But it does not yet give us what we want. It does not provide a way to avoid spinelessness on controversial issues that matter. For one thing, there is no real controversy on any question whose answer is so obvious that it constitutes its own reason, nor when no reason is even required for being justified in answering as one does.

In such simple cases a belief's justification does not derive from its basis in ulterior reasons. Either it derives from a reason provided by the very fact believed, or it may

perhaps derive from one's manifesting a competence in believing as one does, even absent any reason whatsoever. Either way, in such simple cases you would be reasonable to downgrade your opponent based essentially on your disagreement, even with no independent reason for doing so.

This proves insufficient, however, because issues subject to troubling, persistent disagreement are not properly decidable in the absence of ulterior reasons. Controversial issues have no obvious resolutions. They are always complex, and resolvable only through extensive reasoning or evidence. On such an issue it is not so easy to downgrade an opponent based simply on the fact of your disagreement.

4.

Moore's disagreement with the sceptic is an intermediate case. The locus of that disagreement is whether Moore knows about the hand before him. The sceptic

287

says no, Moore says yes. Moore insists by downgrading the sceptic's judgement. True, Moore goes into some detail. But his bottom line is this: the sceptic must be going wrong somehow, since it is more certain that he, Moore, does know about his hand than is anything the sceptic can adduce against that fact.

The sceptic says that he, Moore, might be dreaming. Moore responds that he knows he is not dreaming, if this knowledge is indeed required (as he concedes) for him to know about the hand. His rejection of the sceptic's side on *this* disagreement is thus based on insisting that he does know about the hand. He bases his rejection, in other words, precisely on insisting that he is right in the earlier disagreement and the sceptic wrong. Here he begs the question (whether properly or improperly).<sup>16</sup>

This case differs from any in which one side of a disagreement can be upheld because its correctness is so obvious in itself that no ulterior reason or evidence is required. Moore takes his case definitely *not* to be one where he has no need of ulterior reasons or evidence. He is quite explicit that only based on conclusive reasons does he know himself to be awake and not dreaming. Take again the question whether one can know oneself to be awake. The sceptic answers this question in the negative as against Moore's own affirmative response.

*Moore does not claim it to be obvious that he is awake and not dreaming, so that he needs no ulterior reasons for so believing.* On the contrary, he claims rather that he cannot expound his reasons fully, perhaps because they are too extensive and complex. His reasons are said to constitute 'conclusive evidence' for believing as he does, that he is awake and not dreaming, despite his inability to expound them. Certainly he cannot cite them to an opponent. And Moore might also claim, or so I suggest, that he could not expound them even to himself. Those reasons seem conclusive, then, even if he cannot lay them out, one by one, perhaps even to himself in private.<sup>17</sup>

Moore does not make explicit what he might include in such conclusive evidence. One might however reasonably attribute to him considerations of the sort that seem persuasive to Descartes and perhaps also to J. L. Austin.

288

For Descartes, wakeful experience has a sort of coherence that distinguishes it from dreams; for Austin dreams have a dream-like quality that makes for a similar distinction. Indeed, the distinctions are at bottom the same if what underlies the dream-like quality is just the absence of the coherence that for Descartes distinguishes our wakeful consciousness (often enough, at least when rich enough). Any case of wakeful experience rich enough to have Descartes's coherence and to avoid Austin's dreaminess, would be constituted specifically by some complex stream of consciousness composed of elements that dovetail appropriately with others in the same time slice and also with relevant preceding and succeeding elements. Each such fact of dovetailing presumably contributes to the coherence of the stream of consciousness, and may constitute a 'reason' which, when combined with all the others, provides a conclusive justifying basis for the subject's belief that he is then awake. Each is thus operative as a partial basis for that belief, but Moore apparently believes that this does not require the subject's ability to detail each separately, so as to enable a proof that he could cite to himself or to others.<sup>18</sup>

289

That all being so, Moore does then know, based on evidence in his possession, that he is awake and not dreaming, despite the sceptic's disagreement.<sup>19</sup>

5.

Moore's reasoning suggests a way to handle disagreements commonly encountered in fields where controversy abounds. In Moore's case we are unsure of having fully expounded our evidence.<sup>20</sup> Normal cases of deep, important controversy share relevant features that make this comparison interesting. The evidence on which we base belief in our side of a controversy need only be inscrutable, for whatever reason, or at least sufficiently hard to uncover. For it cannot then be displayed for reflection on how well it supports the content of our belief.

One reason why that body of evidence turns out to be inscrutable recalls Moore's example, where although the evidence is right there in our present conscious experience or remembered directly from our recent experience, *it is too subtle or complex to be noticed and detailed fully*. That is not the main factor, however, in cases of public controversy. Here there is a further source of opacity to our present reflective gaze. Our basis for believing as we do on such questions generally fails to be fully formed and operative in one fell swoop. Light dawns gradually over such questions. A belief forms in us over time through the subtle influence of diverse sources. Some are testimonial, others perceptual, others inferential, and so on. The belief might owe importantly to the believer's upbringing, or to later influence by his community. We are social beings and do well, socially *and* intellectually, to rely on such influence by our social and intellectual

communities. Such proper reliance over time on divergent communities might thus help explain how disagreement can be reasonable.

When our full grounding for a belief is thus complex, and temporally extended with the aid of memory, it lies beyond our present view. Why do we think that our math teacher in high school algebra was named Mrs Hanson? Because we remember that fact. Right, but we must have had reasons for first acquiring and sustaining that belief (which was then held in place through retentive memory). What were those reasons? They are now irretrievable, quite beyond our reach. Why do we take our friend to be touchy on a certain subject. Because on many

290

occasions he has exhibited such a disposition. Yes, that's what we say now when asked for a justification. But we needn't have earlier formed the belief that 'on many occasions Tom has shown himself to be touchy on this subject'. On the contrary, upon experiencing variously and repeatedly how he has responded to relevant promptings, the belief may have dawned on us and guided our future interactions.

Why do we think there are stars in the sky? Can we really cite the reasons that led us to form and retain this belief? Sure, it may be thought, we know we've seen the stars repeatedly; that's our reason. But is that our *only* reason? And, more importantly, what about *this* proposition itself, that we *have* so seen the stars? What is our reason now for believing it? That we ostensibly remember it? We cannot require here the experiential memory of *having seen the stars on a certain occasion*, for this may well be missing. What is said to be enough is rather the ostensible *retentive* remembering of the fact in question. But such a paltry reason is unlikely to exhaust our epistemic justification (even if it does sustain some minimal subjective justification) for believing that there are stars in the sky or that we really have seen them. Even supposing that there is a distinctive 'appearance of retentive memory', which is already dubitable, what is further dubitable is that such an appearance could do much epistemically for our belief in the stars *in the absence* of the relevant past experiences that prompted this belief, a belief then kept in place through the proper operation of retentive memory. It is this latter, more substantial, time- and memory-involving rational basis that need not now be present to our reflective gaze in order to do its proper epistemic work.

And so it is, I submit, for nearly the whole of one's body of beliefs. The idea that we can always or even often spot our operative 'evidence' for examination is a myth.

If we can't even spot our operative evidence—so much of which lies in the past and is no longer operative except indirectly through retained beliefs—then we cannot disclose it, so as to share it. And this will apply to our opinions on complex and controversial topics no less than to our belief about our teacher's name or the touchiness of our friend or the stars in the sky. On all these matters we are in the position that Moore takes himself to be in on the question whether he is awake. We have reasons (in our case through the mediation of retentive memory) that, acting in concert, across time, have motivated our present beliefs,

but we are in no position to detail these reasons fully. This may be so, finally, as Moore also thought, even in cases where the reasons are in combination quite conclusive.

6.

If we leave it at that, however, we risk obscurantism, the no less dispiriting position that our reasons, far removed in our past, or deeply lodged in our

291

subconscious, cannot be uncovered for critical inspection. This can hardly be the way to defend our prized ability to take reasonable stances on controversial topics of politics, religion, philosophy, history, morality, and so forth. We would be trying to defend our ability to take such stances only by removing the topics from the arena of proper debate and critical reflection.

Or so it might seem. But a closer look is revealing. Even if a controversial view that one upholds is justified by much that is then reflectively inaccessible, this by itself need not free one from full responsibility for providing a supporting rationale. We may be justified in insisting on our side of a disagreement despite being unable to spell out our justification. But this does not empower us to proclaim the view unsupported. The demands of public debate are different from those of private belief.

7.

The demands of public debate and public profession require a study of their own. Here we focus rather on reasonable *belief*, which can properly remain unvoiced and undefended. Our emphasis on the importance of hidden reasons can still show that, as in Moore's case, our ability fully to disclose our reasons may be quite limited; which implies in turn that even if reasonable disagreement with full disclosure is impossible, this has little or zero bearing on the possibility of reasonable disagreement, since full disclosure might be difficult and rare.

Nevertheless, let us focus on cases where we do seem able to disclose extensively or even fully. Return first to our two examples used by enemies of reasonable disagreement, that of the contents of the quad, and that of the restaurant bill. In fact such cases can differ broadly. Some might require deference to our opponent, others that we insist on our position, and others that we split the difference. And this may all be so even absent sufficient reason, independently of the substance of our disagreement, to downgrade our opponent's judgement. Cases might still differ dramatically, even if in none can we downgrade independently.

Take the restaurant example. If my friend and I calculate in our heads and come up with contrary results, and we reasonably take each other to be peers on the question at hand (equally good track records, equally good competence as arithmetical calculators, equally good present mental shape, equal attention to the matter, etc.), clearly then each should give equal weight to the other's judgement, so that we ought to 'split the difference' by each moving to a middle position where we both withhold. However, suppose my friend

and I are now out of each other's sight, neither of us privy to what the other does in double-checking our results. I take out pencil and paper, and perform the calculation that way; then I remember the calculator in my pocket, and I perform the calculation that way too. And all of these procedures repeatedly converge on one result: \$42.

292

At the end of the period of rechecking I again compare notes with my friend, who insists that his rechecking has confirmed *his* initial result: \$45. Now things change dramatically. Now I am in the Moore-like position of having to say that if his procedure has led to that result, there must be something wrong with his procedure. I have no idea what it is, since I am not privy to what he did. What has changed from the initial stage is how reliable I can consider my procedure and the competence I manifest in employing it. Nevertheless, I still lack *independent* reason to downgrade my opponent's relevant judgement and his epistemic credentials on the question that divides us. Only based on our disagreement can I now demote him.

The case of the quad may or may not be similar. If the quad is large and the dean barely visible at the opposite end, then my opposing friend's remark can lead me to lower my confidence sharply and with good reason. But if the quad is small and sunlit, and the dean sports a bright blue jacket just fifty feet away below our window, and my friend says there is no one there, and I see the dean there plain as day, then I find it about as plausible to say, Moore-like, that here again something is wrong with my friend, I know not what. And I do this based on the substance of our disagreement.

What makes the difference between the cases where we must yield some ground to our opposing friend, and those where we demote him instead? What makes the difference in such cases is plausibly a matter of how reasonably sure we are of our relevant competence. Take our original restaurant disagreement, or one where the quad is huge and the dean far away. In these cases, because we are sufficiently *unsure* of the competence exercised, the outspoken disagreement of our friend creates a reasonable doubt. When we double-check our addition, however, using paper and pencil and a calculator, and when the conditions of perception are nearly ideal, we reach a very different result. Here we are sure, reasonably so, that the competence we exercise is quite reliable, which is *in part* why we can demote our friend, even if in doing so we must rely on the substance of our disagreement.

That does seem initially plausible about such cases. But a further stance seems more reasonable yet, and compatible with the plausible one just considered. On this further stance the more important difference imported by repeated rechecking is not so much that one is now sure of one's procedure and of its high reliability. The more important difference is rather that one now fails to be sure independently that the opponent's relevant reliability is a match for one's own. So, the important difference is that one no longer has sufficient independent basis for judging the two opponents to be peers on the question at hand.

There is moreover a further difference that seems more important yet. I have been assuming that in the original restaurant case, you and I both (the disagreeing parties) have some degree of confidence in our respective beliefs (\$42 in my case, and \$45 in yours). But we are even more confident that we both are about equally reliable in such mental arithmetic. In the revised restaurant case, there

293

is a crucial change here. After the period of rechecking, when I have checked through paper and pencil and through a calculator, I end up considerably more confident that \$42 is the right result than I am any longer that you are not my inferior on the matter at hand. I have to lower my confidence that you are not inferior because I now know about my greatly increased reliability and I have no such knowledge about you.

There is therefore the following difference between the original case and the revised case. In the original case I am *independently* more sure that you are not inferior than I am that \$42 is the right answer. In the revised case I am more sure that \$42 is the right answer than I am *independently* that you are not inferior. Plausibly, then, our ability reasonably to downgrade our opponent based on the substance of our disagreement varies depending on the degree of confidence we have in our side of the disagreement, compared with the independently based confidence that we have as to whether the opponent is our inferior on the matter at hand.

That is just a matter of internal coherence, of whether we can coherently proceed by downgrading based on the substance of a disagreement. Of course, the *overall* evaluation of whether we proceed correctly in doing so will have to depend also on how reasonably we hold the degrees of confidence that we hold concerning our side of the disagreement, for one thing, and the inferiority of our opponent on the matter at hand, for another thing.

These apparently minor and subtle differences might make a crucial difference to our ability reasonably to sustain our views on controversial issues. For, on such issues our bases for our opinions are not open to view or easily and clearly displayable. Controversial stances are hence liable to differ widely from our stance when we recheck with pencil and paper and a calculator. When we recheck thus our mode of belief acquisition is quite open to view and displayable for all to see. Not so for controversial opinions, or indeed for the great bulk of our standing beliefs now retained through sheer retentive memory. What in the restaurant case is the important difference between the pre-checking and post-checking situations? What is the difference between our pre-checking situation, which calls for suspension, and the post-checking situation, where we may hold our ground and demote the opponent based on the substance of the disagreement? Suppose we locate the important difference in the fact that after rechecking, with pencil and paper and a calculator, we are sure that our method is highly reliable, whereas pre-checking we had used only the unreliable calculation in one's head. This move will put us in dire danger of falling short for any controversial stance that we would like to uphold. After all, most often our position in such stances vis-à-vis our opponents' is not like our position after rechecking in the restaurant case. On the contrary,

we are unlikely to have so much as a clear conception of just how we have formed the opinion that we would like to uphold.

Compare with that the more moderate requirements for proper upholding and for proper downgrading of one's opponent based on the substance of a

294

disagreement. On these more moderate requirements, what relevantly changes in the restaurant case as we move from mere calculation in one's head to elaborate rechecking, is *not* that one is later sure of the reliability of what one can now see clearly as a new and improved basis for sustaining one's side of the disagreement. What relevantly changes is *rather* that one is now no longer independently sure that the opponent is one's peer on the question at hand. And, more relevantly yet, what changes is that after rechecking one is now more sure that one is right on one's side of the disagreement than one is that one's opponent is not inferior on the question under dispute. And this is precisely our position vis-à-vis our opponents on controversial issues, or nearly always so.

8.

What then is the upshot for the possibility of reasonable disagreement on issues of moment in controversial fields? Our reflections inspire cautious optimism. When Pro and Con disagree openly on such an issue, and each insists on his side of the disagreement, each needs to be confident that they proceed competently enough, based on adequate grasp of the available evidence, or at least they need to avoid good independent basis for thinking themselves to be peers on the matter at hand. It will be relatively rare for them to be able to disclose fully their relevant reasons and evidence. To some extent, therefore, they will be relying on a position similar to that of the second restaurant calculator (the one who has rechecked his result). None of them is privy to the backing for their opponents' contrary belief, *not fully*. Yet each might be quite reasonably confident of the competence they themselves exercise, or at least each may have no sufficient independent basis for thinking the other to be a relevant peer. And this is why they might properly downgrade their opponents based essentially on the substance of their disagreement.

Such downgrading need imply no disrespect. Despite your highest respect for the competence of an opponent, you may be forced to downgrade his judgement on a particular issue that divides you. To downgrade him is just to think he must be going wrong *somehow* on that *particular* issue. That this is sometimes appropriately based on the substance of the disagreement, even absent any independent reason to downgrade him, has been a main thesis of this chapter.

If that much is correct, we can then continue to believe as we do on issues of moment in public debate, which does not require us to disrespect our opponents, nor does it exempt us from arguing for our side as fully and convincingly as we can manage. Our inability to defeat an opponent in public debate need not rationally require us to abandon our beliefs. For various powerful reasons, our beliefs can be grounded adequately in reasons that give



us no dialectical advantage, either because they offer no dialectically persuasive leverage, or because they are

295

undisclosably beyond our reach. While appealing to this fact, finally, we must avoid a dispiriting obscurantism, but we need to recognize also that it *is* a fact, one with a clear bearing on our main issue.

We can remain open-minded on *how extensively* the foregoing considerations may help to protect our controversial views. They have some value, we can dare hope, but must be supplemented. For one thing, we must also recognize how often we speak past each other, to one or another extent, when we ostensibly disagree on controversial issues. We need to bear in mind how often this tends to happen, especially in philosophical dialectic.

9.

We have taken note of three factors that may help enable reasonable dissent against open opposition: (a) how verbal disagreement can be, (b) how deeply hidden and undisclosable reasons can be, and (c) how *epistemically effective* a reason can be despite being *dialectically ineffective*.

We have also examined a supposed obstacle to reasonably sustained dissent: namely, that the substance of a disagreement can never provide a proper basis for downgrading one's opponent. We have found this obstacle to be often enough surmountable.

Those four considerations, severally and collectively, open up possibilities for us to sustain our views reasonably even in the teeth of outspoken opposition.<sup>21</sup>

296

#### APPENDIX: MOORE'S DISAGREEMENT WITH THE SCEPTIC

The sceptic and he cannot come together and disclose fully, according to Moore, since he takes his evidence to be inexpressible, or not fully expressible, perhaps opaque even to his own reflection. Nevertheless, Moore's discussion is still relevant in the following way.

There is a broader sense of disagreement where A does not simply affirm the very same proposition that B denies. In this broader kind of disagreement,

A has total evidence  $E(A)$  pertinent to proposition  $P(A)$ , and B has total evidence  $E(B)$  pertinent to proposition  $P(B)$ , where *A would be justified in believing  $P(A)$  based on  $E(A)$  if and only if B would be justified in believing  $P(B)$  based on  $E(B)$* , and where A and B know or believe that fact, or at least the fact that *there is* such evidence on each side.

Moreover: first, A believes that he is and would be justified in believing  $P(A)$  based on his relevant total evidence, whereas, second, B for his part believes that he *is not and would not be justified* in believing  $P(B)$  based on *his* corresponding total evidence.

Moore's is a special case of this.

Moore believes that his sensory experience while wide awake, and his memories of his recent experiences while wide awake, provide him with evidence E(Moore) based on which he is and would be justified in believing Awake(Moore).

Sceptic believes that *his* sensory experience while wide awake, and his memories of *his* recent experiences while wide awake, provide him with total evidence E(Sceptic), based on which *he* is *not* and *would not* be justified in believing Awake(Sceptic).

Finally, they both know or believe that Moore would be justified in believing Awake(Moore) based on his relevant total evidence if and only if Sceptic would be justified in believing Awake(Sceptic) based on *his* corresponding total evidence.

So far Sceptic and Moore do not disagree on a shared question, while sharing also all of their relevant evidence. They do share similar enough evidence, however, evidence that bears similarly enough on relevantly similar propositions in such a way that, as is believed by them both, their respective beliefs in the similar propositions stand or fall together. The sceptic thinks they fall. Moore thinks they stand.

We arrive thus after all at a strict disagreement on one and the same question. On *this* disagreement Moore upholds his side and demotes the disagreeing sceptic by reasoning that makes essential use of a proposition on which they disagree. The proposition that he *does* know about the hand is part of Moore's basis for the conclusion that he must also know himself to be awake and not dreaming, and for the further conclusion that things he knows directly (such as his present and recent experiences) must provide him with conclusive reasons and evidence for his corresponding belief (that he is awake and not dreaming).<sup>22</sup>

## 15 A Justificationist View of Disagreement's Epistemic Significance

Jennifer Lackey

Disagreement with others is a familiar part of our lives. Most of us find ourselves faced with friends who have radically opposing views on the war in Iraq, relatives who explicitly reject our beliefs about religion, and colleagues who dispute our conception of the nature of free will. These are just a few of the more common topics of disagreement, but there are countless others, ranging from the mundane—‘That is a cat by the tree, not a small dog!’—to the extraordinary—‘That I conquered my cancer was not purely a result of my chemotherapy, but was also God answering my prayers.’

In some instances, disagreement can be explained by one party to the dispute being privy to more evidence than the other. You and I may have opposing beliefs about whether Jones committed the murder because only you saw the relevant DNA results linking him to the crime. Other cases of disagreement can be accounted for in terms of various kinds of cognitive asymmetries. You and I may disagree about whether the bird in the backyard tree is a starling because only I am using eyeglasses with an out-of-date prescription. Still other instances of disagreement can be explained by one member of the dispute having reason to believe that the other is epistemically inferior in some respect. You and I may have differing attitudes about whether my daughter is the best ballet dancer in her class because you believe that I am biased where my daughter's talents are concerned.<sup>1</sup> But adjusting our doxastic states in all of these sorts of cases does not reveal anything significant about disagreement itself, since the fact that you and I disagree drops out of the explanation of this adjustment; it is, for instance, the difference in our familiarity with the relevant evidence or the asymmetry in our cognitive capacities that does the explanatory work. In order to truly assess the significance of *disagreement* itself, there cannot be any relevant epistemic asymmetries between the parties to the dispute to shoulder the explanatory

298

burden;<sup>2</sup> in other words, such parties should be epistemic peers.<sup>3</sup> The question at issue, then, is this: what is the significance of disagreement between those who are epistemic peers? In particular, what is the rational response to disagreement in situations where there are no relevant epistemic asymmetries between the members involved in the dispute?

There are two answers to this question found in the recent literature. On the one hand, there are those who hold that disagreement itself can be wholly without epistemic significance; thus, one can continue to rationally believe that  $p$  despite the fact that one's epistemic peer explicitly believes that not- $p$ , even when one does not have a reason independent of the disagreement in question to prefer one's own belief. I shall call those

who hold this view *nonconformists*. According to nonconformists, there can be reasonable disagreement among epistemic peers. So, for instance, Gideon Rosen writes:

It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable. (Rosen 2001: 71)

Similarly, Thomas Kelly claims that:

The mere fact that others whom I acknowledge to be my equal with respect to intelligence, thoughtfulness, and acquaintance with the relevant data disagree with me about some issue does not undermine the rationality of my maintaining my own view.

(Kelly 2005: 192)<sup>4</sup>

According to nonconformists, then, the mere fact that there is disagreement with one's epistemic peer does not mandate any sort of doxastic revision from either party to the dispute.

Now, there are two central explanations of the nonconformist response to peer disagreement.<sup>5</sup> First, there is what we may call the *egocentric view*. On this view, I am justified in giving my belief extra weight<sup>6</sup> in the face of peer disagreement because the belief in question is *mine*. So, for instance, Ralph Wedgwood claims:

Perhaps, quite generally, it is rational for one to place greater trust in one's own intuitions, simply because these intuitions are one's own, than in the intuitions of other people. In other words, perhaps it is rational for each of us to have an *egocentric epistemic bias* in favour of our own intuitions. (Wedgwood 2007: 261)<sup>7</sup>

299

According to the egocentric version of nonconformism, then, the epistemic symmetry involved in peer disagreement can be broken by virtue of the extra weight afforded to my belief in virtue of the belief's being mine.<sup>8</sup>

The second explanation of the nonconformist response to peer disagreement is what we may call the *correct reasoning view*. On this view, I am justified in giving my belief extra weight in the face of peer disagreement because the belief in question is *in fact the product of correct reasoning*. Thus, Thomas Kelly writes:

The rationality of the parties engaged in [peer disagreement] will typically depend on who has in fact correctly evaluated the available evidence and who has not. If you and I have access to the same body of evidence but draw different conclusions, which one of us is being more reasonable (if either) will typically depend on which of the different conclusions (if either) is in fact better supported by that body of evidence.

(Kelly 2005: 180)

So, according to the correct reasoning version of nonconformism, peer disagreement's epistemic symmetry can be broken by virtue of the extra weight afforded to my belief in virtue of its being in fact best supported by the evidence.<sup>9</sup>

In contrast to nonconformists, there are those who hold that disagreement itself possesses enormous epistemic significance; thus, unless one has a reason that is independent of the disagreement itself to prefer one's own belief, one cannot continue to rationally believe that  $p$  when one is faced with an epistemic peer who explicitly believes that  $\text{not-}p$ . The requirement that the reason for preferring one's own belief be 'independent of the disagreement' at issue is meant to include both independence from the belief in question and from the reasoning grounding this belief. Without this requirement, begging the question would presumably be an appropriate way to respond to a case of peer disagreement—for example, one could settle a dispute over the question whether  $p$  by appealing to  $p$  itself. Moreover, since any instance of disagreement will provide one disputant with a reason to believe that the other has failed at some point in her reasoning, parties to a dispute would never qualify as epistemic peers if appealing to the reasoning grounding the belief in question were permissible. I shall call those who hold this view *conformists*.<sup>10</sup> According to conformists, there cannot be reasonable disagreement among epistemic peers. Thus, Richard Feldman claims that:

in situations of full disclosure, where there are not evident asymmetries, the parties to the disagreement would be reasonable in suspending judgement about the matter at hand.

300

There are, in other words, no reasonable disagreements after full disclosure, and thus no mutually recognized reasonable disagreements. (Feldman 2006: 235)<sup>11</sup>

In a similar spirit, David Christensen says that:

in [cases of disagreement with epistemic peers], I should change my degree of confidence significantly toward that of my friend (and, similarly, she should change hers toward mine). (Christensen 2007: 189)<sup>12</sup>

And Adam Elga writes:

When you count an advisor as an epistemic peer, you should give her conclusion the same weight as your own . . . call it the 'equal weight view'. (Elga 2007: 478)

Suppose that before evaluating a claim, you think that you and your friend are equally likely to evaluate it correctly. When you find out that your friend disagrees with your verdict, how likely should you think it that you are correct? The equal weight view says: 50%. (Elga 2007: 488)<sup>13</sup>

Conformists, then, argue that equal weight should be given to one's own beliefs and to those held by one's epistemic peers, and thus significant doxastic revision is required in the face of peer disagreement. What kind of doxastic revision is necessary? Answers to this question vary. Feldman, for instance, casts the debate in terms of an all-or-nothing model of belief, and so he argues that disagreement with an epistemic peer regarding the

question whether  $p$  requires that both parties to the dispute *withhold* belief relative to  $p$ . Christensen and Elga instead frame the issues in terms of degree of belief, and so they argue that disagreement with an epistemic peer regarding the question whether  $p$  requires *splitting the difference* in the degrees of their respective beliefs. Thus, where 1 represents maximal confidence that  $p$  is true and 0 represents maximal confidence that  $p$  is false, if I give credence 1 to the proposition that Smith committed the murder and you give credence 0 to the proposition that Smith committed the murder, our attitude towards  $p$  should converge in the middle—we should give credence .5 to this proposition and become perfect agnostics with respect to Smith's guilt. But regardless of the details, conformists all agree that when epistemic peers disagree, substantial adjustment is required in their respective beliefs.

301

Despite the differences between nonconformists and conformists, they appear to share a commitment to a thesis that I shall call *Uniformity*,<sup>14</sup> which can be characterized as follows:

*Uniformity*: Disagreement with epistemic peers functions the same epistemically in all circumstances.

According to this thesis, it doesn't matter whether one's beliefs conflict with an epistemic peer's over a confidently held perceptual experience or a dubious political conclusion, a necessary mathematical proof or a supernatural religious doctrine, simple directions to the store or a complicated philosophical view—disagreement with epistemic peers either always does or does not require doxastic adjustment.

But what, one might ask, is involved in being an epistemic peer with someone? Answers to this question typically involve requiring the satisfaction of at least the following two conditions:

*Evidential equality*:<sup>15</sup> A and B are evidential equals relative to the question whether  $p$  when A and B are equally familiar with the evidence and arguments that bear on the question whether  $p$ .

*Cognitive equality*:<sup>16</sup> A and B are cognitive equals relative to the question whether  $p$  when A and B are equally competent, intelligent, and fair-minded in their assessment of the evidence and arguments that bear on the question whether  $p$ .<sup>17</sup>

302

In addition to evidential and cognitive equality, Richard Feldman adds what he calls full disclosure to the conditions relevant to the disagreement at issue. More precisely:

*Full disclosure*: A and B are in a situation of full disclosure relative to the question whether  $p$  when A and B have knowingly shared with one another all of their relevant evidence and arguments that bear on the question whether  $p$ .<sup>18</sup>

Let us say that when there is both evidential and cognitive equality between A and B in situations of full disclosure with respect to the question whether  $p$ , they are *epistemic peers*.

I shall call disagreement involving epistemic peers in this sense *idealized*, which can be understood as follows:

*Idealized disagreement*: A and B disagree in an idealized sense if and only if, relative to the question whether  $p$ , (1) A and B are aware that they hold differing doxastic attitudes, (2) prior to recognizing that this is so, A and B take themselves to be epistemic peers with respect to this question,<sup>19</sup> and (3) A and B are epistemic peers.

The ‘aware that’ clause of condition (1) is included to rule out the relevance of the following sort of case: I believe that  $p$  and some person in China, who happens to be my epistemic peer relative to this question, believes that not- $p$ , but we are entirely unaware both of each other and of our disagreement. Surely, it is not even clear that such a case properly involves a disagreement, let alone one that should be at the centre of this discussion. Similarly, condition (2) is included to preclude the relevance of the following sort of case: I believe the rather complicated mathematical conclusion that  $p$  and a three-year-old, whom I

303

have just met, believes that not- $p$ . While I am aware of our disagreement, I have absolutely no idea that this three-year-old is a math whiz for her age and thus my epistemic peer regarding this question. Once again, such a case clearly does not represent the sort of disagreement at issue in this debate.

Now, idealized disagreement is to be distinguished from what I shall call *ordinary disagreement*. Ordinary disagreement does not require the parties to the dispute to *actually* satisfy the conditions of evidential equality, cognitive equality, and full disclosure. In particular:

*Ordinary disagreement*: A and B disagree in an ordinary sense if and only if, relative to the question whether  $p$ , (1) A and B are aware that they hold differing doxastic attitudes, and (2) prior to recognizing that this is so, A and B take themselves to be roughly epistemic peers with respect to this question.<sup>20</sup>

So, whereas idealized disagreement occurs when the parties to the dispute are, as a matter of fact, epistemic peers, ordinary disagreement takes place when such parties at least take themselves to be roughly such peers.<sup>21</sup> There are, then, two different questions that may be at issue: first, can one continue to rationally hold a belief in the face of idealized disagreement and, second, can one continue to rationally hold a belief in the face of ordinary disagreement? Which of these is the focus of the debate between nonconformists and conformists?

Given that the distinction between idealized and ordinary disagreement is not drawn in the literature, combined with the fact that some theorists seem to emphasize the former

while others rely on the latter, it is not entirely clear that there is a single answer to this question. For instance, Peter van Inwagen's emphasis on evaluating disagreement in terms of 'all objective and external criteria' (van Inwagen 1996: 275) and Elga's requiring that peers 'have *the same* evidence' (Elga 2010: ms p. 2) suggest idealized disagreement, while Feldman's talk of there being no '*evident* asymmetries' (Feldman 2006: 235) between the parties to the debate indicates that he is concerned with ordinary disagreement. And some theorists appear to focus on both kinds of disagreement, but at different points in their discussion. For instance, Kelly motivates the problem of epistemic disagreement with the question, 'Can one rationally hold a belief while knowing that that belief is not shared (and indeed, is explicitly rejected) by individuals over whom one possesses no discernible epistemic advantage' (Kelly 2005: 168), which sounds very similar to what I am calling ordinary disagreement. But when he explicitly argues against conformists, Kelly requires

304

evidential and cognitive equality between the parties to the debate, which parallels what I am calling idealized disagreement.<sup>22</sup> So, while both kinds of disagreement figure in the debate, they do not always do so explicitly or consistently. In what follows, I shall, when relevant, be clear that there is the distinction between idealized and ordinary disagreement and I shall specify which kind is at issue when necessary.

In this chapter, I shall argue that neither nonconformism nor conformism provides a plausible account of the epistemic significance of peer disagreement. I shall first show that in some cases, nonconformism provides the intuitively correct result, and in other cases, conformism does. This leads to the rejection of Uniformity: disagreement with epistemic peers does not function epistemically the same in all circumstances. I shall then develop my *justificationist* account of peer disagreement's epistemic significance. Whereas current views maintain that disagreement, by itself, either simply does or does not possess epistemic power, my account holds that its epistemic power, or lack thereof, is explainable in terms of the degree of justified confidence with which the belief in question is held. I shall then show that my justificationist account has two central advantages: first, it is able to provide a principled explanation for why nonconformism provides the intuitively correct result in some cases, while conformism gives the intuitively correct result in other cases and, second, it is generalizable in a way that neither of these rival views is.

## 1. NONCONFORMISM

There are two questions that are at the centre of the debate between nonconformists and conformists: (1) does disagreement with an epistemic peer require substantial doxastic revision, and (2) can there be reasonable disagreement between epistemic peers? As we saw above, an answer to the former is taken to dictate an answer to the latter:<sup>23</sup> nonconformists respond negatively to (1) and thus affirmatively to (2),<sup>24</sup> while conformists answer affirmatively to (1) and thus negatively to (2).<sup>25</sup> In this section, I shall



begin with (1), focusing primarily on the response nonconformists have given to this question.

305

After doing so, I shall return to a more general consideration of both of these questions.

To my mind, the most promising line of defence for the nonconformist begins with the following type of case:

PERCEPTION: Estelle, Edwin, and I, who have been room-mates for the past eight years, were eating lunch together at the dining room table in our apartment. When I asked Edwin to pass the wine to Estelle, he replied, 'Estelle isn't here today'. Prior to this disagreement, neither Edwin nor I had any reason to think that the other is evidentially or cognitively deficient in any way, and we both sincerely avowed our respective conflicting beliefs.

Now, PERCEPTION can be read as involving either idealized disagreement or ordinary disagreement. Let us evaluate these in turn.

If Edwin and I are in an idealized disagreement over the presence of Estelle, then we must be epistemic peers with respect to this question, which requires evidential and cognitive equality in a situation of full disclosure. But it is unclear how to make sense of disagreement occurring in PERCEPTION under these conditions. For recall that two people are evidential equals relative to a question when they are equally familiar with the evidence and arguments that bear on this question. However, if I have a phenomenologically vivid experience of it seeming to me that Estelle is at the dining room table, and Edwin does not, then how could we be equally familiar with the evidence that bears on whether Estelle is present? Perhaps evidential equality could be glossed as follows: two people are evidential equals when they have an equal amount of evidence supporting their given beliefs. So, Edwin and I need not be equally familiar with the same relevant evidence; we just need to possess equal amounts of evidence for our respective, conflicting beliefs. If this weaker notion of evidential equality is granted, however, then pressure begins to build against the plausibility of the cognitive equality condition obtaining. For at least one of us must be hallucinating or experiencing some sort of cognitive malfunction in order to plausibly explain how one of us claims to see Estelle while the other does not when we are all presumably inches from each other. It is, therefore, quite difficult to even grasp how there could be idealized disagreement between Edwin and me in this situation. For these reasons, PERCEPTION seems best understood as a case of ordinary disagreement.

Let us, then, regard the disagreement between Edwin and me as ordinary in nature. Now, consider the situation in question from my perspective: it clearly seems to me that Estelle is sitting at the dining room table with me—indeed, suppose that minutes earlier we were engrossed in conversation while eating our pasta. Moreover, I have never in my life hallucinated an object, I have not been drinking or taking any drugs, I have my contact lenses in, I have ample evidence of my eyesight functioning reliably when my

nearsightedness is corrected, and I know all of this to be true of myself. How, then, should I rationally respond to

306

Edwin's claim that Estelle is not dining with us? Despite the fact that, up to now, I have had good reason to regard Edwin as an epistemic peer, it seems clearly rational for me to continue to believe just as strongly that Estelle is present at the table. Indeed, even after full disclosure—where Edwin explains that he does not seem to see anything in the chair that Estelle purportedly occupies—I still seem rational in being fully convinced that she is dining with us at the table. For given the extraordinarily high degree of justified confidence with which I hold my belief about Estelle's presence, Edwin's disagreement seems best taken as evidence that something has gone awry with him, either evidentially or cognitively. In other words, I seem justified in concluding that Edwin is no longer an epistemic peer, even if he was prior to the disagreement in question. This conclusion is further evidenced by considering whether it would be appropriate for my belief to continue to guide my responses and actions in ways paradigmatic of confident belief. Imagine, for instance, that immediately after Edwin's full disclosure of his reasons for disagreeing with me, the doorbell rings and I open the door to find Estelle's mother asking if she is at home. Surely, it would be rational for me to respond affirmatively to this question. Indeed, Estelle's mother would rightly be utterly perplexed if I were to say, 'Although I can apparently see her, I really do not know if she is here since I am withholding belief in light of my disagreement with Edwin on the topic'. Thus, when the disagreement in question is ordinary, nonconformism seems to deliver the appropriate intuitive response in PERCEPTION.<sup>26</sup>

307

Similar considerations apply in the following case:

ELEMENTARY MATH: Harry and I, who have been colleagues for the past six years, were drinking coffee at Starbucks and trying to determine how many people from our department will be attending the upcoming APA. I, reasoning aloud, say, 'Well, Mark and Mary are going on Wednesday, and Sam and Stacey are going on Thursday, and, since  $2 + 2 = 4$ , there will be four other members of our department at that conference'. In response, Harry asserts, 'But  $2 + 2$  does not equal 4'. Prior to this disagreement, neither Harry nor I had any reason to think that the other is evidentially or cognitively deficient in any way, and we both sincerely avowed our respective conflicting beliefs.

As was the case in PERCEPTION, there seem to be two general points that can be made about the situation described in ELEMENTARY MATH. First, when this case is said to involve idealized disagreement, it becomes rather inexplicable. How could two adults—both of whom are functioning normally—who possess equal evidence relevant to the question at hand, disagree on whether  $2 + 2 = 4$ ? Surely, at least one of us is either confused or cognitively deficient in some way. Second, when the disagreement is ordinary, it intuitively seems quite rational for me to retain my belief even in the face of

disagreement with Harry, whom I have had very good reason to believe is an epistemic peer. For given my extraordinarily high degree of justified confidence in my belief that  $2 + 2 = 4$ , Harry's disagreement seems rightly regarded by me as evidence that he is not well, either evidentially or cognitively. In other words, as was the case in PERCEPTION, Harry's disagreement with me over the truth of  $2 + 2$  equalling 4 seems appropriately taken by me as evidence that we are no longer epistemic peers, and thus nonconformism again provides the intuitively plausible response.

Of course, even when the disagreement in PERCEPTION and ELEMENTARY MATH is ordinary, rather than idealized, it may be argued that the kind of disputes in these cases is so peculiar that it is unclear whether any general conclusions about nonconformism follow from them. For instance, how often does it happen that people disagree about whether their friend is sitting in the chair next to them, or about whether  $2 + 2 = 4$ ? So, let us consider a case that is similarly mundane, but where the ordinary disagreement in question is more likely to obtain:

DIRECTIONS: I have lived in Chicago for the past fifteen years and during this time I have become quite familiar with the downtown area. Of the many restaurants that I enjoy frequently dining at, My Thai on Michigan Avenue is among my favourites. Jack, my neighbour, moved into the same apartment building the very weekend that I did fifteen years ago and he, too, has become quite competent in his acquaintance with the city. Indeed, it is not uncommon for us to bump into each other at various places, My Thai being one of them. Today, when I saw Jack coming out of his apartment, I told

308

him that I was on my way to My Thai on Michigan Avenue, after which he responded, 'My Thai is not on Michigan Avenue—it is on State Street'. Prior to this disagreement, neither Jack nor I had any reason to suspect that the other's memory is deficient in any way, and we both rightly regarded one another as peers as far as knowledge of Chicago is concerned.

What response should I rationally have to Jack's ordinary disagreement with me about My Thai's location? In particular, must I withhold, or at least significantly reduce my confidence in, my belief because a neighbour whom I believe to be an epistemic peer claims that the restaurant is on State Street? To my mind, nonconformism once again seems to give the correct intuitive result. For if I have lived in Chicago for fifteen years, know the city extremely well, frequently eat at My Thai, have not been drinking or taking any drugs, have substantial evidence that my memory is functioning reliably, and know all of this to be true of myself, then I seem perfectly justified in my confidence about My Thai's location, even after Jack fully discloses his vivid memory of the restaurant being on State Street. Indeed, given the substantial amount of credence and epistemic support enjoyed by my belief, Jack's disagreement seems appropriately regarded as evidence that something is not right with him. I may, for instance, suspect that he has been drinking, is delusional, or is suffering from some kind of memory loss; in any case, I seem rational in concluding that Jack is no longer an epistemic peer regarding the location of My Thai.

Thus, as was the case in PERCEPTION and ELEMENTARY MATH when the disagreement in question is ordinary, nonconformism seems to deliver the intuitively correct result in DIRECTIONS.<sup>27</sup>

But, the conformist may object, in order for you to rationally retain your fully confident belief in the face of disagreement with an epistemic peer, there needs to be what we may call a ‘symmetry breaker’.<sup>28</sup> A symmetry breaker is something that indicates that the epistemic position of one of the parties to the disagreement in question is superior to the other's. For instance, in DIRECTIONS, I have had good reason for believing that Jack is my epistemic peer relative to My Thai's location, and thus I have had good reason for believing that our epistemic situations are symmetrical regarding this topic. In order for me to fully retain my confidence in my belief in the face of disagreement with Jack, then, there needs to be a symmetry breaker that justifies my nonconformity. Otherwise, my resistance to doxastic revision is little more than dogmatic egoism.

Here is where the ordinary nature of the disagreement is relevant. For notice: even when you and I have had excellent reasons for believing that we are epistemic peers, I will, in ordinary situations, often have access to information

309

about myself that I lack with respect to you—let us call this *personal information*. Personal information is information that one has about the normal functioning of one's own cognitive faculties. I may, for instance, know about myself that I am not currently suffering from depression, or not experiencing side effects from prescribed medication, or not exhausted, whereas I may not know that all of this is true of you. Of course, given that I have excellent reason for believing that you and I are epistemic peers, I cannot, apart from the disagreement itself, have evidence for believing that you *are* currently suffering from depression, or that you *are* experiencing side effects from prescribed medication, and so on. But this is certainly compatible with my not knowing that these things are not true of you. Consider: two oncologists, both of whom are widely regarded as the best in the field, can surely appropriately regard one another as epistemic peers relative to cancer diagnoses, even if they do not at any given time know that the other is not depressed, not overtired, not suffering from a personal crisis, and so on. Now, applying these considerations to DIRECTIONS, note that the very disagreement at issue indicates that at least one of us—that is, either Jack or myself—is seriously cognitively malfunctioning. However, I know about myself that I have not been drinking, have not suffered from any recent delusions, and do not have any evidence for questioning the reliability of my memory. And while I do not have prior reason to question Jack's capacities, the fact of the disagreement itself, in conjunction with the personal information that I possess about myself, now gives me reason to think that there *is* a serious problem with his cognitive faculties. In particular, my personal information, when it combines with the already extraordinarily high degree of justified confidence that I have in my belief in DIRECTIONS, is able to serve as the requisite symmetry breaker here. Where Jack had previously seemed to be an epistemic peer of mine, I no longer

regard him as one—his epistemic standing has collapsed. This permits me, then, to retain my belief about the restaurant's location with the same high degree of justified confidence I had before the disagreement occurred.

This same point could be cast in the language of defeaters. An instance of ordinary disagreement regarding the question whether  $p$  provides me with a defeater for my belief that  $p$ . When I am very highly justified in holding this belief, the personal information that I possess about myself and lack about my interlocutor can provide me with a defeater-defeater for this belief. And, so long as I do not then acquire a defeater-defeater-defeater, I am thereby permitted to rationally retain my belief that  $p$  with the same degree of credence.

Perhaps the conformist will here object that the above considerations reveal precisely why it should be idealized disagreement that is at issue. For stipulating that the parties to the disagreement are, in fact, epistemic peers who have fully disclosed their equal evidence effectively rules out the sort of symmetry breaker provided by personal information in cases where there is a high degree of justified confidence, thereby enabling us to focus entirely on the epistemic significance of the disagreement itself. In particular, it can be built directly into evidential and

310

cognitive equality that both parties to the debate know of each other that there are no relevant asymmetries in their respective epistemic situations.

Notice, first, however, that when the case is idealized to this extent, it becomes quite difficult to make sense out of how disagreement is even possible. For if everything *even remotely relevant* to the topic at hand must be equal, then epistemic peers begin to sound much more like epistemic *clones*. It then becomes perplexing how epistemic clones relative to a question can even be engaged in a disagreement. Second, even if we were able to render coherent idealized disagreement between epistemic peers in DIRECTIONS, this concept has virtually no connection to the very disagreements that breathe life into this debate. For it is very common for philosophers writing on this topic to motivate interest in it by citing debates in history, philosophy, politics, religion, and other areas where disagreement is widespread and impassioned. But these debates bear very little resemblance to the hyper-idealized scenarios under consideration here. It would rarely, if ever, happen that two people continue to disagree with one another about, say, gay marriage where there is evidential and cognitive equality with known epistemic symmetry in situations of full disclosure.<sup>29</sup> Typically, there are all sorts of asymmetries at work, such as in the background assumptions that are being tacitly relied upon—for example, only one party to the dispute may regard what the Bible says as relevant to social institutions like marriage—or in different character traits—for example, only one party may be risk-averse—or in different values—for example, only one party may value tradition more than equality, and so on. Moreover, it is even rarer for the two parties to the debate to know of one another that there are no epistemic asymmetries, particularly when personal information is at issue. How often does it happen, for instance, that I know

that my colleague, with whom I disagree about the Iraq war, is not depressed, exhausted, distracted, and so on, on any given day? Thus, conclusions drawn from hyper-idealized situations involving disagreement ultimately tend to have very little connection to the disagreements we face every day.<sup>30</sup> For these reasons, I think that, at least for the most part, ordinary disagreement ought to be the focus of this debate.<sup>31</sup>

There is another door that is opened for the nonconformist by focusing on ordinary disagreement. Recall that Feldman adds the condition of full disclosure to the kinds of disagreement that are at issue in this debate. But obviously, there are more and less idealized versions of this requirement. At one end of the spectrum, full disclosure requires that epistemic peers knowingly share with one another literally all of the evidence and arguments in their possession that bear in any way on the question under dispute. While no doxastic revision is

311

hard to rationally justify under such hyper-idealized circumstances, it is also difficult to understand how or whether this type of disagreement ever in fact obtains. As Ernest Sosa has emphasized, many of our beliefs are supported by countless pieces of subtle and complex evidence acquired via multiple sources over a number of years.<sup>32</sup> When disagreeing over deeply entrenched religious and political beliefs or over memorial beliefs from the distant past, for instance, it is often practically impossible to conjure up all of the evidence and arguments that we have that bear on these topics. Once again, then, idealization worries lead to focusing on ordinary disagreement, which requires only that the parties to the debate take themselves to be roughly epistemic peers on the topic at hand. This, however, opens the door to relevant epistemic asymmetries explaining the disagreements at issue. For if there are subtle and complex pieces of evidence that have not been fully disclosed, then some of these very pieces of evidence may explain why the epistemic peers are engaged in the dispute under consideration. Perhaps underwriting my belief is a phenomenological experience that I cannot adequately convey, or massive amounts of evidence accumulated over many years that I couldn't possibly remember, or data acquired from various sources and contexts that I am unable to articulate. In these types of cases, disclosure would, at best, be only partial. Given that these sorts of epistemic asymmetries may be at work in cases of ordinary disagreement, then, there may be rational reasons for epistemic peers to not engage in doxastic revision.

So, when we focus on ordinary disagreement, it looks as though nonconformists are correct that a negative response to question (1) above is warranted: a subject who disagrees with an epistemic peer is *not* thereby rationally required to revise her beliefs. Does this thereby necessitate, as nonconformists suggest, an affirmative answer to question (2)—that is, that there can be reasonable disagreement between epistemic peers? Not necessarily.<sup>33</sup> For while all three cases above are ones where I am not required to revise my doxastic attitudes in the face of disagreement, they also involve my coming to conclude that something has gone awry—either evidentially or cognitively—with my companion. For instance, upon hearing that Edwin is sincerely denying that Estelle is

dining with us in PERCEPTION, I no longer believe that we are evidential and cognitive equals with respect to this question. Thus, the disagreement that Edwin and I have regarding this question is not a reasonable one; that is to say, it is not the case that we are equally reasonable in holding our respective beliefs.

But then is it correct to say that the nonconformist's negative response to question (1) has been defended? For if I no longer regard you as evidentially and cognitively equal in the above cases, then has it been shown that doxastic revision is not rationally required in the face of disagreement *between epistemic peers*? In

312

order to answer this question, recall Elga's characterization of conformism in the following passage:

Suppose that before evaluating a claim, you think that you and your friend are equally likely to evaluate it correctly. When you find out that your friend disagrees with your verdict, how likely should you think it that you are correct? The equal weight view says: 50%. (Elga 2007: 488)

According to conformism, then, if pre-disagreement (at  $t_1$ ) A believes that she and B are equally likely to be correct regarding the question whether  $p$ , then post-disagreement (at  $t_2$ ) A should believe the same. In other words, the disagreement itself should not change one's beliefs about the probability that one is right. Yet this is precisely what I, along with the nonconformist, deny in the above cases. For in PERCEPTION, ELEMENTARY MATH, and DIRECTIONS, I believe at  $t_1$  that in the case of disagreement with my epistemic peer, the probability that I am right regarding the question whether  $p$  is 50% and then at  $t_2$ , in light of the nature of the disagreement itself and the positive personal information that I possess, I believe that the probability that I am right is dramatically higher.<sup>34</sup> Given this, it would be more accurate to frame the original questions with which we began as follows: (1\*) does disagreement with someone whom, were it not for the disagreement in question, one would regard as an epistemic peer, require substantial doxastic revision, and (2\*) can there be reasonable disagreement between those whom, were it not for the disagreement in question, would regard one another as epistemic peers? It should now be clear that while I have defended the negative answer given by nonconformists to (1\*), I do not thereby endorse their positive answer to (2\*). That is to say, disagreement with someone who, were it not for the disagreement in question, one would regard as an epistemic peer, does not necessarily require doxastic revision, but it does not follow from this that there is *reasonable* disagreement.

It may be argued, however, that the disagreement in question is not actually changing one's beliefs about the probability that one is right in the above cases. For instance, consider my belief about the probability that I am right in my belief

313

about Estelle's presence in PERCEPTION. If you were to ask me what probability I would assign to my being right were I to disagree with Edwin on some unspecified

perceptual issue, I would say 50%. But if you were to ask me what probability I would assign to my being right were I to disagree with him regarding, say, whether Estelle is sitting inches away from us, wouldn't I assign a much higher probability to my being right?<sup>35</sup> Doesn't this, then, show that I am not downgrading the epistemic status of my peer on the basis of the disagreement itself, which is precisely what is at issue between the nonconformist and the conformist? For wouldn't the probability that I assign to my being right on the question whether Estelle is sitting inches away from us be the same at both  $t_1$  and  $t_2$ ?

No, and here is why. If I genuinely regard Edwin as my epistemic peer where perceptual matters are concerned, then even when asked in the abstract what probability I would assign to my being right were we to disagree regarding whether Estelle is inches from us, I would say 50%. It is only in the context of the *actual disagreement itself*—where I have a phenomenologically vivid experience of Estelle sitting inches from me and have relevant personal information about the normal functioning of my cognitive faculties—that I would offer a significantly lower probability. Given this, I downgrade Edwin's status as my peer precisely because of the disagreement in question, and thus the probability that I assign to my being right regarding whether Estelle is sitting inches from us differs between  $t_1$  and  $t_2$ . This conclusion is clearly at odds with the view of disagreement advanced by the conformist.

At this point, then, there are three conclusions that I have reached: first, there are some cases where idealizing the conditions at issue renders the relevant disagreement either inexplicable or disconnected from the very disagreements that motivate the debate; second, there are some cases of ordinary disagreement where nonconformism delivers the intuitively correct result, at least with respect to the question whether I must revise my beliefs in the face of this disagreement; and, third, personal information can combine with the already high degree of justification held by a confident belief so as to provide precisely the symmetry breaker that is needed to explain why nonconformism seems to get it right in these cases.<sup>36</sup> More will be said about these conclusions, but let us now consider whether any intuitive support can be garnered for conformism.

314

## 2. CONFORMISM

One of the more intuitively compelling cases on behalf of conformism is provided by Christensen, which can be presented as follows:<sup>37</sup>

**BILL CALCULATION:** While dining with four of my friends, we all agree to leave a 20% tip and to evenly split the cost of the bill. My friend, Ramona, and I rightly regard one another as peers where calculations are concerned—we frequently dine together and consistently arrive at the same figure when dividing up the amount owed. After the bill arrives and we each have a clear look at it, I assert with confidence that I have carefully calculated in my head that we each owe \$43 and Ramona asserts with the same degree of confidence that she has carefully calculated in her head that we each owe \$45.<sup>38</sup>



Now, the first point to notice is that, as we saw with the earlier cases, hyper-idealizing the conditions in BILL CALCULATION renders the disagreement in question rather inexplicable. For instance, if evidential *equality* is required, it may be argued that I have the experience of going through a calculation that seems to support one-fifth of the bill being \$43, and Ramona has the experience of going through a calculation that seems to support one-fifth of the bill being \$45, so it is unclear how our evidence could be equal. Perhaps what is meant is merely that we have an equal *amount* of evidence supporting our different beliefs? But this surely cannot be what evidential equality requires, for there will then be countless cases where peer disagreement is explainable in terms of one party to the dispute being privy to a piece of crucial evidence that the other lacks. Moreover, full disclosure of the disagreement at issue will presumably include our sharing our calculations with one another, at least one of which is incorrect. So, after this full disclosure, what explains how idealized disagreement between epistemic peers persists in such a case?

Given these considerations, BILL CALCULATION, like the earlier cases, seems best understood as involving ordinary disagreement. Thus, apart from the

315

disagreement at hand, Ramona and I both take ourselves to be epistemic peers with respect to the question under consideration—that is, we take ourselves to be in roughly a position of both evidential and cognitive equality regarding the amount each of us owes and to have fully disclosed our evidence to one another.<sup>39</sup> What, then, is the rationally appropriate response to our disagreement? Christensen writes, ‘it seems quite clear that I should lower my confidence that my share is \$43 and raise my confidence that it's \$45. In fact, I think (though this is perhaps less obvious) that I should now accord these two hypotheses roughly equal credence’ (Christensen 2007: 193). While Christensen argues that Ramona and I should split the difference in the degrees of confidence in our respective beliefs, Feldman's view is that each of us should withhold belief about the amount owed; in both cases, however, the conformist requires substantial doxastic adjustment in the face of ordinary disagreement with peers.

I must admit to sharing the conformist's intuitions in BILL CALCULATION. Given that I argued on the side of nonconformism with respect to PERCEPTION, ELEMENTARY MATH, and DIRECTIONS, but admit that conformism provides the correct result here, I now want to consider whether there is a principled explanation of this difference.

The first point to notice about the two types of cases is the degree of confidence with which the beliefs in question are held. In all of the cases supporting nonconformism, I am extremely confident in the truth of the beliefs that I hold—I am, for instance, extraordinarily sure that Estelle is sitting next to me at the table, that  $2 + 2 = 4$ , and that My Thai is on Michigan Avenue. In an ordinary case of disagreement, it would take a great deal more for me to adjust my doxastic states with respect to these beliefs than one person disagreeing with me, even one for whom I had, until that moment, good reason to believe was an epistemic peer. In contrast, when I divide in my head the amount of a

large bill owed by five of us, I may be confident in my calculation, but I never come anywhere near the confidence that I have in my belief that a friend is currently sitting inches from me, or that  $2 + 2 = 4$ , or that a restaurant that I have been frequently going to for the past fifteen years is on a street as unforgettable as Michigan Avenue.

316

To test the relevance of the degree of confidence with which beliefs are held, consider the following:

**MODIFIED BILL CALCULATION:** While dining with four of my friends, we all agree to leave a 20% tip and to evenly split the cost of the bill. My friend, Ramona, and I rightly regard one another as peers where calculations are concerned—we frequently dine together and consistently arrive at the same figure when dividing up the amount owed—and we both have a clear view of the bill. After repeatedly going through the division of the amount of the bill on paper, I assert with a high degree of confidence that we each owe \$43. Ramona, who was also busy performing the division on paper, asserts with the same degree of confidence that we each owe \$45.

While some doxastic revision may still be appropriate here, I do not at all have the same conformist intuitions that I have in **BILL CALCULATION**. And I suspect that if the case is modified to include more and more rounds of calculation done by me on paper, or perhaps even on a calculator, we may ultimately end up altogether eliciting nonconformist intuitions, even if Ramona is busy doing the same number of calculations. For in a wide range of cases, the amount of doxastic revision rationally required by ordinary disagreement seems to diminish as the degree of confidence with which the belief in question is held increases.

But, clearly confidence, even when it is extraordinary, cannot be the only relevant feature here. For this would have the consequence that the hyper-dogmatist—who is supremely confident in all of her beliefs—is never rationally required to adjust her doxastic states in the face of ordinary disagreement with epistemic peers. This brings us to the second point about the difference between those cases that support nonconformism and those that support conformism: what **PERCEPTION**, **ELEMENTARY MATH**, and **DIRECTIONS** have that **BILL CALCULATION** lacks is a highly confident belief that is *very well-justified epistemically*. In the first three cases, I am in optimal epistemic conditions relative to the beliefs in question: I see my room-mate of eight years sitting inches away from me in excellent lighting, I fully grasp with cognitive faculties that are not deficient in any way the truth of  $2 + 2 = 4$ , and I vividly remember My Thai being on Michigan Avenue, where this is the most famous street in Chicago and the restaurant is one that I have walked to countless times in the past fifteen years. Given these optimal epistemic conditions, the probability of my being wrong in any of these three beliefs is extremely low, and thus my very confident beliefs are clearly highly justified.

In **BILL CALCULATION**, however, the situation is quite different: I arrive at a belief after dividing in my head the amount of a large bill owed by five of us. Several features

of this case render the likelihood of my being wrong somewhat high. For instance, the bill is rather large—if it were instead a \$25 bill that I was dividing equally by five, Ramona asserting that we each owe \$6 rather than \$5 would fail to clearly elicit conformist intuitions. I am also dividing the bill by five—if I were dividing a \$218 bill by two, my believing that we each owe \$109 may not intuitively require doxastic adjustment in the

317

face of an epistemic peer's disagreement. Moreover, I am doing the calculation in my head—if I were working the math out on paper, we may be less inclined to say that I should withhold belief given Ramona's disagreement with me. Finally, I perform the calculation only once—if I had repeatedly done the math and consistently arrived at the belief that we each owe \$43, conformism no longer seems to be the clear response to the situation. What these considerations reveal is that the likelihood of my being wrong in BILL CALCULATION is significant, and thus neither my level of confidence nor my level of justification is very high with respect to the belief in question. To my mind, the absence of a highly justified confident belief is what explains the clear conformist intuitions in this case.

This provides the resources for answering a question that some readers may have at this point: why is there a symmetry breaker in the cases supporting nonconformism, but not in BILL CALCULATION? Here is the answer: personal information, when it combines with the already high degree of justification possessed by a confident belief—such as that enjoyed by my belief that  $2 + 2 = 4$ —is sufficient for breaking the prior epistemic symmetry between you and me when we disagree. But the cases that clearly elicit conformist intuitions can be different from those that intuitively justify nonconformism in various ways. First, the nature of the disagreement in the latter cases indicates that at least one party to the dispute is seriously cognitively malfunctioning. When there is a disagreement regarding whether  $2 + 2 = 4$  or a friend is sitting at the table, for instance, it is unlikely that this can be explained by appealing to ordinary errors. Second, as noted above, the confidence enjoyed by the beliefs in the cases supporting conformism may not be very high. In BILL CALCULATION, for instance, it is plausible to think that I regard it as quite possible that I am wrong. Third, even if I tend to think very well of my abilities and thereby have an unusually high degree of confidence in my belief, this high degree is surely not justified given the substantial fallibility in this situation. Given these differences, the positive support provided by one's personal information will not be adequate, when combined with the relatively low degree of justified confidence, to render ordinary disagreement with a peer epistemically benign in cases supporting nonconformism.

### 3. CONSEQUENCES

We are now in a position to draw some conclusions about the epistemic significance of disagreement. First, the cases where nonconformism clearly provides the correct result are ones where there is a symmetry breaker between one's epistemic peer and oneself that

is provided by the presence of personal information combining with a highly justified confident belief. More precisely:

318

**No Doxastic Revision Required:** In a case of ordinary disagreement between A and B, if A's belief that  $p$  enjoys a very high degree of justified confidence<sup>40</sup>, then A is permitted to rationally retain her same degree of belief that  $p$  if and only if A has a relevant symmetry breaker.

Second, the cases where conformism clearly provides the correct result are ones where there is a relatively low degree of justified confidence such that the positive support provided by personal information is insufficient for breaking the epistemic symmetry between one and one's epistemic peer. In particular:

**Substantial Doxastic Revision Required:** In a case of ordinary disagreement between A and B, if A's belief that  $p$  enjoys a relatively low degree of justified confidence, then A is rationally required to substantially revise the degree to which she holds her belief that  $p$ .

There will, then, be many cases that fall on the spectrum between no doxastic revision required, and substantial doxastic revision being necessary. If, say, A's belief that  $p$  enjoys a moderately high degree of justified confidence, then merely some doxastic revision may be required in the face of ordinary disagreement with an epistemic peer. For instance, rather than withholding belief or splitting the difference in degree of belief, A may be required to only somewhat reduce the degree to which she believes that  $p$ . On the other hand, if A's belief that  $p$  enjoys a moderately low degree of justified confidence, then more doxastic revision may be required in the face of ordinary disagreement, but perhaps still not as much as withholding or splitting the difference in degree of belief.<sup>41</sup>

Third, in spite of my siding with nonconformists in some cases and conformists in others, it is a mistake to regard my view as merely a hybrid of the two. I do not agree with the nonconformist that peer disagreement is entirely without epistemic power, even when I support the intuitive results of such a view; nor do I agree with the conformist that peer disagreement possesses substantial epistemic power, even when I support the intuitive results of this view. In this sense, I reject both nonconformism—because the absence of doxastic revision in the face of peer disagreement is never justified merely by virtue of the fact that beliefs are either mine or are the product of correct reasoning—and conformism—because substantial doxastic revision in the face of peer disagreement is never justified merely by virtue of equal weight being given to my own beliefs and to those held by my epistemic peers. Instead, my view holds that peer disagreement's epistemic power, or lack thereof, depends on the degree of justified confidence with which

319

the belief in question is held combined with the presence or absence of relevant personal information. It is for this reason that my view is *justificationist* rather than simply a blend of nonconformism and conformism.

What notion of justification is at work here? While a detailed response to this question lies beyond the scope of this chapter, I should say at least a few words about how I am understanding this concept. To this end, let us briefly return to PERCEPTION which, it may be recalled, involves Edwin and me disagreeing over whether Estelle is present at the dining room table in our apartment. Now, let us suppose that Edwin's denial that Estelle is present is the result of hallucinating which, in turn, is caused by the fact that Edwin has been unknowingly drugged by a friend. The drug in question produces no signs discernible to the subject of its effects, and Edwin has no independent reason to doubt his own cognitive capacities.<sup>42</sup> While I am happy to grant that, from a purely subjective point of view, Edwin is just as reasonable in his belief regarding Estelle's presence as I am in mine, I would not agree that our beliefs are even close to being equally justified. In particular, I am enough of an externalist about justification to require that the process or faculty responsible for the production of the belief in question be reliable or otherwise appropriately truth-conducive.<sup>43</sup> Thus, in the face of disagreement with one another, Edwin and I would both be entitled to hold our conflicting beliefs with the same degree of credence only to the extent that such beliefs are produced by processes that are equally reliable or truth-conducive.<sup>44</sup> Given that one belief is the result of a veridical perceptual experience, and the other the result of a hallucination, my view clearly grants radically different justificatory statuses to them.

It should also be noted that the level of justified confidence operative in the antecedents of the NDRR and the SDRR is determined prior to the disagreement in question, but the disagreement itself can affect whether there is ultimately a relevant symmetry breaker. This point can be illustrated by using the language of defeaters. An instance of peer disagreement regarding the question whether  $p$  may be interpreted as providing me with a defeater for my belief that  $p$ . However, if my belief that  $p$  has a very high degree of justified confidence, then the personal information that I possess about myself and lack about my interlocutor can become salient and thereby provide me with a defeater-defeater for my original belief. This then enables me to rationally retain my belief that  $p$  after the disagreement with the same degree of credence. If, however, my belief that  $p$  does not have a high degree of justified confidence prior to the disagreement, then the force of the disagreement can do the defeating work in question, thereby requiring substantial doxastic revision on my part.

Fourth, as should be clear, I reject Uniformity. My view maintains that ordinary disagreement with a peer sometimes permits nonconformist results and, at other times, conformist results, depending on whether a symmetry breaker is provided in conjunction with the degree of justified confidence enjoyed by the belief in question. I take this to be a significant virtue of the present account, as it not only accommodates the apparently conflicting intuitions in the literature, but also provides an explanation of them.

Finally, unlike responses offered by conformists to counterexamples such as PERCEPTION and DIRECTIONS, my view provides an explanation of why no doxastic

revision is required in such cases that is both principled and plausible.<sup>45</sup> To see this, consider the following slightly modified case from Christensen (2007):

**EXTREME BILL CALCULATION:** While dining with four of my friends, we all agree to leave a 20% tip and to evenly split the cost of the bill. My friend, Mia, and I rightly regard one another as peers where calculations are concerned—we frequently dine together and consistently arrive at the same figure when dividing up the amount owed. After the bill arrives and we each have a clear look at it, I assert with confidence that I have carefully calculated in my head that we each owe \$43. In response, Mia asserts with the same degree of confidence that she has carefully calculated in her head that we each owe \$450, which is more than the total cost of the bill.

By way of response to this sort of case, Elga offers the following:

according to the equal weight view, your probability that you are right should equal your prior probability that you would be right, *conditional on what you later learn about the circumstances of the disagreement*. And one circumstance of the split-the-check disagreement is that you are *extremely* confident that your advisor's answer is wrong—much more confident than you are that your answer is right. Indeed, her answer strikes you as obviously insane. So in order to apply the equal weight view, we must determine your prior probability that you would be right, conditional on these circumstances arising.

To do so, think of your state of mind before doing the calculation. We have assumed that, conditional on the two of you disagreeing, you think that your advisor is just as likely as you to be right. But it is also natural to assume that, conditional on the two of you disagreeing *and your finding her answer utterly insane*, you think that you are much more likely to be right. If so, then when that circumstance arises the equal weight view instructs you to favor your own answer. That is the intuitively correct verdict about the case.

What makes the above answer work is an asymmetry in the case. You find your advisor's answer insane . . . . (Elga 2007: 491)<sup>46</sup>

321

So, according to Elga, conformists need not appeal to justification in order to accommodate the intuition that no doxastic revision is required in such a case, for there is a symmetry breaker in EXTREME BILL CALCULATION that can do the work all on its own; namely, that I find Mia's answer utterly insane. Given this, similar remarks could perhaps be extended to cover the counterexamples from Section 1, such as PERCEPTION and DIRECTIONS—I possess a symmetry breaker in both cases because I find the responses offered by Edwin and Jack utterly insane.

There are, however, several problems with this response. First, if the mere fact that one finds the view of one's opponent utterly insane suffices to provide a symmetry breaker in cases of peer disagreement, then not only is the very spirit of the conformist position put in serious jeopardy, but the door is also opened for bias and dogmatism to justify one in rarely engaging in doxastic revision. For isn't it the case that the most impassioned sexists

typically find arguments on behalf of women's rights laughable, the most devout racists find the views of proponents of equality crazy, and the most committed homophobes find the idea of granting homosexuals even the most basic dignities perverse. But even putting aside such extreme versions of bias and dogmatism, isn't it quite common for those engaged in deep, impassioned debates to find the views of their opponents completely insane? Opponents of abortion, for instance, frequently argue that the position of pro-choice advocates is incomprehensible, those against the war in Iraq are often at a complete loss trying to grasp the view of its supporters, and it is not at all uncommon for critics of Andy Warhol's art to have difficulty understanding the minds of its admirers. Yet given what Elga tells us in the passage quoted above, it looks as though doxastic revision would not be required in such cases since finding the view of one's opponent 'utterly insane' provides the necessary symmetry breaker. Surely, this is a most unwelcome result for an account such as Elga's that pitches itself as an *equal weight* view.

Once the conformist attempts to avoid this problem, however, it becomes clear that all of the work supporting nonconformist intuitions cannot be done without appealing to the notion of justified confidence, which brings us to the second problem with Elga's response: the possession of symmetry breakers makes sense in the cases under consideration only against the background of beliefs that have a very high degree of justified confidence. To see this, notice that in the cases at issue—PERCEPTION, DIRECTIONS, and EXTREME BILL CALCULATION—the beliefs opposing mine are not, in and of themselves, utterly insane.<sup>47</sup> By itself, it is not utterly insane for Edwin to believe that Estelle

322

is not at the lunch table with us—this is insane only against the background of my very highly justified belief that Estelle is sitting inches away from us at the lunch table. In and of itself, it is not evidently insane for Jack to believe that My Thai is on State Street—this is extremely doubtful only when viewed against my very highly justified belief that it is on Michigan Avenue. And viewed in isolation, it is not evidently insane for my friend to believe that we each owe \$450 when we are splitting our restaurant bill evenly among five of us—this is insane only against the background of my extremely highly justified belief that the amount owed by each of us when we evenly split the bill must be lower than the total bill. Thus, we can see that the symmetry breakers that Elga claims are doing all of the work are in fact efficacious only when there are highly justified beliefs underwriting them, thereby providing further support for my justificationist account of disagreement.

Third, as suggested above, even if utter insanity could be understood without smuggling a high degree of justification through the back door, this response on behalf of the conformist would fail to generalize to include all of the problematic cases. For there is nothing about Jack's response in DIRECTIONS that is utterly insane in any reasonable sense. Moreover, it can certainly be stipulated that in all of the relevant cases, both parties

to the disagreement find one another's answer equally insane. Given this, there wouldn't even be the possibility of the 'utter insanity' symmetry breaker that Elga has in mind, and yet doxastic revision is still not intuitively required on my part in such cases.

Similar considerations apply to Christensen's response (2007) to cases such as PERCEPTION, DIRECTIONS, and EXTREME BILL CALCULATION. According to Christensen, I do not have to engage in doxastic revision in the face of peer disagreement so long as the best explanation I can offer of the disagreement lies with an error made by my interlocutor rather than with one made by me. For instance, in EXTREME BILL CALCULATION, Christensen claims that 'The real ground for thinking that my friend made the error . . . derives from the fact I have evidence that my assessment of the disputed proposition is supported by an extremely reliable kind of reasoning, but I have no basis for supposing the same about my friend's contrary assessment' (2007: 20). But there are two readings of this. On a subjective reading, so long as the best explanation *from my subjective point view* of the disagreement in question lies with an error made by my peer, no doxastic revision is required. This reading, however, opens the door to rampant bias and dogmatism justifying an extreme version of nonconformism. For so long as one is egocentric or biased enough, one can always convince oneself that the best explanation of a given instance of peer disagreement lies

323

with errors made by one's interlocutors. In BILL CALCULATION, for instance, I could be so dogmatic about my calculations that it just seems obvious to me that the best explanation of my disagreement with Ramona lies with reasoning errors on her part. This is obviously an unwelcome result. Given this, an objective reading of Christensen's clause may be more promising, according to which no doxastic revision is needed so long as the best explanation *from an epistemic point of view* of the disagreement in question lies with an error made by my interlocutor. But this reading leads to my justificationist account of disagreement's epistemic significance. For instance, what makes Edwin's error the best epistemic explanation of our disagreeing over Estelle's presence is the extraordinary degree of justified confidence I have in my belief that she is sitting inches from us at the table, which thereby enables the personal information that I have about myself to serve as a symmetry breaker. If both this level of justified confidence and personal information were removed, it would be utterly mysterious what would make Edwin's error the best epistemic explanation of our disagreement over Estelle's presence. Thus, on this reading of Christensen's clause, the best epistemic explanation of the disagreement in question just bottoms out in a combination of justified confidence and the presence or absence of relevant personal information.

There is one final problem with the conformist response that I should like to mention, one that afflicts the views of both Elga and Christensen. Recall that conformists explicitly require that the reason for preferring one's own belief in a peer dispute be independent of the disagreement in question, where this is meant to include both independence from the belief and from the reasoning grounding this belief. Both of the claims offered above,



however, are at odds with this requirement. With respect to Elga's response, the charge of utter insanity, as mentioned above, makes sense only when features of the actual disagreement are taken into account. For instance, Edwin's denial of Estelle's presence in the room is insane to me only against the background of the high degree of justified confidence possessed by my belief that she is so present. But then breaking the symmetry between Edwin and me by appealing to the utter insanity of his answer *is not independent of the grounds for my belief*. Regarding Christensen's response, I know that my belief is the result of an extremely reliable process precisely because of the nature of the grounds of this belief, for example, my vivid phenomenological experience of Estelle. Once again, then, breaking the symmetry between Edwin and me is *not independent of the grounds for my belief*. Thus, in an attempt to accommodate cases of extreme disagreement, conformists undermine a central thesis of their view.

Unlike rival views, then, my justificationist account has the resources to explain why no doxastic revision is required in cases such as PERCEPTION, DIRECTIONS, and EXTREME BILL CALCULATION in both a principled and plausible fashion while appealing to just the original resources of the view.

324

#### 4. CONCLUSION

In this chapter, I have defended a justificationist view of the epistemic significance of ordinary disagreement among epistemic peers. Such an account provides a plausible explanation for why nonconformism delivers the intuitively correct result in some cases of peer disagreement while conformism provides the right response in others. These are significant virtues that the present view has over its rivals, as it reveals that my justificationist view is a fully generalizable account of disagreement's epistemic significance that provides principled and unifying explanations of intuitions that would otherwise appear to be in conflict.<sup>48</sup>

325

## Notes

### 1. INTRODUCTION

<sup>1</sup> It is really preferable to view preservationist and expansionist conceptions as representing *degrees* of distance or departure from traditional epistemology. There are no sharp dividing lines here. But often it is useful to talk in terms of binary or discrete categories, even when the underlying reality is messier than such categorization suggests.

<sup>2</sup> An exception is Steve Fuller, who used ‘social epistemology’ both as the title of his first book (Fuller 1988) and as the title of a journal he founded. Another revisionist, Martin Kusch (2002), uses a somewhat similar phrase for his position: ‘communitarian epistemology’.

<sup>3</sup> Whether it succeeds is another question. I have argued that the activities described by Latour can be reinterpreted in respectably epistemic terms (Goldman 1999: 227–30).

<sup>4</sup> I am not prepared to endorse this denial in my own voice. For example, can't one be a legitimate epistemologist and still endorse some type of epistemic relativism about justification or rationality? And what about a philosopher who examines arguments for and against relativism without taking a firm position of her own? Isn't she doing real epistemology—despite the fact that one of the ‘core’ assumptions of traditional epistemology is the falsity of epistemic relativism? My purpose here, however, is only to explain or reconstruct the probable reasoning of those epistemologists who, according to Alston, would refuse to accept social epistemology as ‘real’ epistemology. I am not necessarily joining them in excommunicating (even) rejectionist social epistemology. Of course, rejectionist social epistemology isn't the kind I personally favour. But it's a more delicate question to decide whether it deserves to be considered epistemology at all.

<sup>5</sup> Another formulation Boghossian offers of the key thesis of classical relativism is a thesis he calls ‘Equal Validity’: ‘There are many radically different [mutually incompatible] yet “equally valid” ways of knowing the world, with science being just one of them’ (Boghossian 2006: 2). Again, a claim of equal validity goes beyond any semantical thesis of New Age relativism.

<sup>6</sup> In addition, there are paradigmatic questions of epistemology to which social *answers* have been offered—even if the questions themselves do not invoke social properties or relations. As long as these answers are defended in a characteristically epistemological fashion, this should also be considered traditional—albeit social—epistemology.

<sup>7</sup> The term ‘evidence’ is used here in two different but related senses. First, it can refer to an object or state of affairs in the external world. This is the sense in which we speak of

rings in a tree trunk being evidence (for the age of the tree) and of fossils like Lucy and Ida being evidence (about the evolution of Homo Sapiens). In addition to this external sense, there is a second sense of ‘evidence’ in which it refers to belief, justified belief, or knowledge *about* evidence in the external sense. Thus, to possess, gather, or acquire evidence is to have or acquire beliefs, justified beliefs, or knowledge about external evidence. In this chapter, ‘social evidence’ is sometimes used in the external sense, referring, for example, to statements uttered or written by other people, or their opinions (states of mind). Other occurrences of ‘social evidence’ refer to subjective states of a doxastic agent concerning external social evidence. I trust that the intended use will be clear from the context.

<sup>8</sup> Two volumes dedicated to these questions are Lackey and Sosa (2006) and Feldman and Warfield (2010). Most of the articles in the present volume also address these questions.

<sup>9</sup> See Goldman and Shaked (1991). A proof of the relevant theorem appears in the reprinted version of the paper. The proof presupposes a linear measure of truth possession, which many would question because it isn't a ‘proper scoring rule’. Nonetheless, other approaches to epistemic utility might offer related rationales for preferring one test to others (Maher 1990; Fallis 2007). Any such rationale would suffice for present (merely illustrative) purposes, so the argument does not depend essentially on the assumptions used in the Goldman–Shaked analysis.

<sup>10</sup> See Unger (1975), Williamson (1996, 2000), DeRose (2002), and Hawthorne (2004).

<sup>11</sup> Why doesn't Robert Brandom's (2000) discursive or communicational theory appear on this list? Because Brandom's targets of analysis are *semantical* concepts, not *epistemic* ones. He is not, fundamentally, trying to illuminate epistemic justifiedness.

<sup>12</sup> The literature on dialectical, or dialogical, approaches to logic and argumentation is substantial. It ranges from treatments oriented toward formal logic to treatments attuned to informal logic, natural conversation, and norms of argumentation. Prominent specimens of this work include Hamblin (1970), Barth and Krabbe (1982), Lorenzen (2001), Toulmin (1958), and Woods et al. (2000).

<sup>13</sup> Alston doesn't suggest that any epistemologists hold precisely this position. He says that many would regard ‘much’ of the material in *Knowledge in a Social World*—not all of it—as not real epistemology.

<sup>14</sup> Someone might argue that what makes irrationality interesting for traditional epistemology is the concern with doxastic decision-making processes of individual agents. But collective agents don't really engage in doxastic decision-making, one might claim. So why should a branch of epistemology take any interest in their irrationality? I dispute the assumption that collective agents do not engage in doxastic decision-making (DDM). Although collective agents are (presumably) not distinct centres of consciousness, this does not mean that they don't engage in DDM. ‘Doxastic decision-making’ refers to whatever processes transpire to bring about, in a subject, a new or

revised doxastic state. Which process this is in the case of collective agents is up for grabs. It might just consist in a set of distributed processes in their constituent members, together with an aggregation process that works automatically, without any ‘extra’ individual activity. An assumption that needs to be avoided is that group DDM must involve things like consciousness, attention, or purposefulness *in* the group entity.

<sup>15</sup> In *Knowledge in a Social World* I restricted the relevant epistemic outcomes to ‘veritistic’ outcomes, i.e., true belief, false belief, and intermediate degrees of belief associated with the truth. In general what I mean by ‘epistemic outcomes’ in this paper is a wider family of epistemically valued and disvalued states, including justified and unjustified belief, and so forth. In the legal context, however, I am again interested in evaluations in terms of ‘veritistic’ outcomes.

<sup>16</sup> Section 4.2 may have suggested that doxastic decision-makers must always take the initiative in gathering evidence. But as we see in the legal trial example, social systems can play a crucial role in producing or assembling social evidence relevant to the doxastic task at hand. In legal trials, the court issues subpoenas to relevant witnesses and requires them to appear in court and testify, which involves answering questions posed by attorneys. Witnesses’ testimony, of course, is evidence for use by the jury. The jury itself, which is the doxastic agent, does not have to collect such evidence—indeed, they are barred from collecting evidence external to the trial situation. Many other social systems serve a roughly similar evidence-assembling function. Traditional enterprises like entire educational systems and public libraries are intended to do exactly this sort of thing. In the modern era the Internet and its assorted applications serve this function. A salient illustration is the online encyclopedia, Wikipedia. This may be viewed as a project in mass collaboration that uses certain methods to produce a large collection of social evidence (writings produced by multiple authors). How reliable and trustworthy is this evidence? This is another suitable topic for expansionist social epistemology (see Fallis and Sanger 2009).

<sup>17</sup> For further elaboration of this point, see Goldman (1999: 298–300).

<sup>18</sup> For defence of the thesis that many epistemological assessments in mainstream individual epistemology are fundamentally based on reliability considerations (or other truth-linked considerations), see Goldman (1979, 1986).

<sup>19</sup> Let me briefly elaborate this point. Can impartiality be understood as not giving one party *any* advantage over the other in resolving the dispute? This is not quite right. If one party to the dispute has a dozen witnesses on her side of the story of what occurred whereas the opposing party has only two, the system confers an ‘advantage’ on the former by allowing all witnesses to testify. Does this mean that the system is unfair or partial? Certainly not. Does impartiality consist in giving both parties to a dispute an equal chance of winning? This would suggest that flipping a coin would be a good decision procedure; but that would be absurd. A fair or impartial system, then, is better understood as a system that confers no advantage to a party *except* in virtue of her merit-determining properties, i.e. properties in virtue of which she deserves to win. If being on

the side of the *truth* (as concerns the material facts of the case) is a principal merit-determining property, then producing decisions in accordance with the truth is what makes a system fair and impartial. According with the truth is, of course, a primary epistemic virtue (Goldman 1999: 290–1).

<sup>20</sup> This is how I have classified it in two earlier treatments of the topic, Goldman and Cox (1996) and Goldman (1999: ch. 7). In both cases it was argued that the strong form of the free-market-for-speech thesis, in which it is defended as a consequence of economic theory, is false. However, the specific position I take on the subject is not at issue here. The only issue here is whether this qualifies as a legitimate question of *epistemology*.

<sup>21</sup> Franz Dietrich (2007) argues, however, that there is no notion of probability such that the conditions of competence and independence are likely to be satisfied simultaneously.

<sup>22</sup> Another example is Michael Fuerstein (2009), who emphasizes the role of scientific method in democratically appropriate deliberation.

<sup>23</sup> However, we shouldn't preclude a classification of 'theoretical' for possible work of this kind. The Condorcet Jury Theorem is plausibly classified as a theoretical contribution to social epistemology. Another example—directly germane to the topic of group deliberation—is a formal result by Christian List (forthcoming), which doesn't, however, bear on issues of reliability.

<sup>24</sup> I refer to Larry Laudan's *Truth, Error, and Criminal Law* (2006), reviewed by Raphael Goldman and Alvin Goldman (2009) in the journal *Legal Theory*.

<sup>25</sup> Thanks to Blake Roeber, whose comments on the penultimate draft led to several helpful clarifications. Earlier drafts of this material were presented at the Stirling conference on social epistemology, a University of Calgary graduate student conference, and the World Congress of Philosophy in Seoul, Korea (2008). I thank the audiences in all these venues for stimulating discussion and criticism.

## **2 Testimony, Advocacy, Ignorance: Thinking Ecologically about Social Knowledge**

<sup>1</sup> I discuss this example more fully in Code (1995).

<sup>2</sup> Moran's view is largely consistent with Michael Welbourne's position in chs. 5 and 6 of *Knowledge*, and in his *The Community of Knowledge*. See Welbourne (1986, 2001).

<sup>3</sup> See Code (1987).

<sup>4</sup> See Sosa (1991), Steup (2004), and Goldman (2008).

<sup>5</sup> See in this regard Bergin (2002).

<sup>6</sup> Wittgenstein observes, 'Knowledge is in the end based on acknowledgement' (1968: §378).

<sup>7</sup> See Welbourne (1979: 1–9, 2001: 109–13).

<sup>8</sup> I discuss Fricker (2007) in my (2008).

<sup>9</sup> See Castoriadis (1991, 1994, 1998).

<sup>10</sup> Code (2006).

<sup>11</sup> The phrase is from Donna Haraway (1991).

<sup>12</sup> Here I draw on Anderson (1995).

<sup>13</sup> I elaborate this resistance in ch. 3 of my (2006).

<sup>14</sup> Latour asks: ‘When will we be able not to reduce matters of concern . . . to matters of fact?’ (2004: 51).

<sup>15</sup> See Harding (1993).

<sup>16</sup> See Kroll (2002), Lytle (2007), and Smith (2001).

<sup>17</sup> IDRC stands for the International Development Research Centre. I discuss this case more fully in Code (2008).

<sup>18</sup> From an unidentified letter-writer to the *New Yorker*, reprinted in its 1995 anniversary issue (cited in Smith 2001: 741).

<sup>19</sup> For Fricker, hermeneutical injustice exposes ‘a lacuna in the collective hermeneutical resources’ (2007: 150). Hence, a woman unable to claim a hearing for her suffering from sexual harassment, before the concept became available, and a man unable to own his ‘nascent identity as a homosexual’ in 1950s America (163–5), exemplify the harms hermeneutical marginalization inflicts on a person's or a group's sense of self.

<sup>20</sup> IDRC (International Development Research Centre) Online Publications: ‘FIXING HEALTH SYSTEMS’, 2004; [http://www.idrc.ca/en/ev-64404-201-1-DO\\_TOPIC.html](http://www.idrc.ca/en/ev-64404-201-1-DO_TOPIC.html) ‘Refreshed’: 14 August 2007.

<sup>21</sup> ‘Before TEHIP [Tanzania Essential Health Interventions Project] we did not identify and prioritize our interventions,’ recalls Dr Harun Machibya, Morogoro District Medical Officer. ‘Rather, we implemented plans worked out centrally. Even in budgeting, the tendency was to add some percentages to previous years' planned and budgeted activities.’

<sup>22</sup> See Sackett et al. (2000) and Goodman (2003).

<sup>23</sup> See especially ch. 5, ‘Patterns of Autonomy, Acknowledgement, and Advocacy’.

<sup>24</sup> Recall Anne Seller: ‘As an isolated individual, I often do not know what my experiences are’ (1988: 180).

<sup>25</sup> Bowden (2008: 52).

<sup>26</sup> Cynthia Cockburn remarks, in another context, ‘the researched take greater risks, and lack the mobility, resources and choices available to the researcher’ (1998: 3).

<sup>27</sup> See Lear (1977: 433–52).

<sup>28</sup> See e.g. Lytle (2007: Epilogue and Afterword)

<sup>29</sup> See *An Inconvenient Truth*, a film directed by Davis Guggenheim; and Al Gore (Guggenheim and Gore, 2006).

<sup>30</sup> ‘Subjugated knowledges’ is a term from Michel Foucault (1980: 82).

<sup>31</sup> See Proctor and Schiebinger (2008), and Sullivan and Tuana (2006, 2007).

<sup>32</sup> Williams (1997) exposes the harms of such neutrality. See also Mills (2008). Mills observes: ‘The point of trying to understand white ignorance is, of course, *normative* and not merely sociological: the goal of trying to reduce or eliminate it’ (2008: 235).

<sup>33</sup> Letter from Clarence Cottam to Carson, February 1962, quoted in Lear (1977: 401).

### **3 Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time**

<sup>1</sup> For an early case for the view that justification can be dispersed in the scientific community, see John Hardwig (1985: 335–49). Lynn Hankinson-Nelson (1990) has argued for the view that the scientific community is the subject of scientific knowledge.

<sup>2</sup> Hardwig (1985).

<sup>3</sup> See Fricker (2007).

<sup>4</sup> Williams's paper was originally delivered as part of a Symposium on ‘Expanding Epistemology’, with Guy Axtell and Robert Brandom, APA Pacific Division Meeting, Spring 2007. The present chapter grew out of my response as Commentator on that occasion.

<sup>5</sup> Note that Brandom himself is ready to go a step further in the externalist direction than Williams is willing to. In *Articulating Reasons* he embraces reliabilism's ‘Founding Insight’ and allows that, for instance, an expert in distinguishing Toltec from Aztec potsherds can know whether a shard is one or the other even if she cannot say how she does it. Williams differs, and so maintains a stronger internalism in his responsibilist position. See Brandom (2000: ch. 3, esp. 98–9).

<sup>6</sup> This is the characterization of Cartesian scepticism that Williams has elsewhere argued to be the only one to furnish radical scepticism—the view that we have no justified beliefs. See, for instance, Williams (2001: 73–7).

<sup>7</sup> This is how Brandom (2000b) puts the issue.

<sup>8</sup> This distinction was first made by Bernard Williams, as Craig notes. See Williams (1973: 146).

<sup>9</sup> This is Michael Welbourne's term for it. See his (2001), especially ch. 6.

<sup>10</sup> Here I paraphrase somewhat but intend to capture Craig's conditions. For his own formulation, see Craig (1990: 85).

<sup>11</sup> Properties such as these bestow what Bernard Williams (2002: 42–3) in his genealogy of truthfulness (modelled closely on Craig's genealogy of knowledge) calls 'purely positional advantage'.

<sup>12</sup> To be precise, there are three key pressures that push the good informant's proto-knowledge towards the objectivized form it takes as, simply, knowledge. First, sometimes enquirers may not need to recognize any informant here and now, but only at some point in the future. Second, the enquirer may be aware that there can be good informants whose indicator properties he is unable to detect. And third, it may not matter to the informant that he himself acquires the information at all, as what may matter is simply that someone around here has got it. All three push the idea of knowledge in the direction of 'objectivization' and away from any dependence on immediate subjective availability to the enquirer.

<sup>13</sup> See Craig (1990: 62–3).

<sup>14</sup> In *Unnatural Doubts* (1991) Williams argues for contextualism independently from considerations about Default and Challenge, since it pre-dates the publication of Brandom's *Making It Explicit* (1994)—I thank Alessandra Tanesini for first pointing this out to me. In Williams's later book, *Problems of Knowledge* (2001), we see the two in combination.

<sup>15</sup> Most directly of all, our genealogical story undermines Cartesian scepticism characterized in the traditional manner as a demand for absolute certainty, for that demand is most immediately what the proto-contextualism generated in the State of Nature exposes as fatally misguided, in as much as there is no practical context where a good informant would have to be absolutely certain in the requisite sense of indubitability. But the present focus is on the more potent interpretation of the Cartesian challenge as threatening radical scepticism.

<sup>16</sup> Bernard Williams (2002) exploits this facility by identifying in the State of Nature his two fundamental virtues of truth—Accuracy and Sincerity—and then going on to explore their contingent historical forms and significances.

<sup>17</sup> 'Epistemological context' already seems too generous a category, precisely because there are already approaches (such as genealogical approaches) to epistemology that preempt, or at least do not invite, sceptical challenge.

<sup>18</sup> An earlier draft of this chapter was presented at a conference on Social Epistemology held at the University of Stirling in 2007, and I am grateful to all those present for their input, including Kathleen Lennon, who responded. I also thank Duncan Pritchard for subsequent written comments. The present chapter is a slightly revised version of an article that appears in *Philosophical Papers*. For details see the bibliographical entry for Fricker (2008). I am grateful to the editor of *Philosophical Papers* for permission to reprint.



#### **4 On Saying that Someone Knows: Themes from Craig**

<sup>1</sup> Some well known exceptions and criticisms of this are found in Bishop and Trout (2005a, 2005b), Goldman (1992a), Kornblith (2002), Stich (1990), and Weinberg, Nichols, and Stich (2001).

<sup>2</sup> See Fricker (1998, 2007), Neta (2006), Lane (1999), and Kusch (2009).

<sup>3</sup> See Kaplan (2003).

<sup>4</sup> Cf. the careful discussion of similar themes in Rysiew (2007).

<sup>5</sup> This chapter has been long in the making. Thanks to Erik Olsson, Jesper Kallestrup, and Mikkel Gerken for very helpful and stimulating discussions along the way. The material was presented at ‘The Place of Epistemic Agents’ conference, Madrid, 2–3 October 2008. I would like to thank the organizers for providing me with the opportunity to present the material at that occasion, and the audience for helpful suggestions.

#### **5 The Swamping Problem Redux: Pith and Gist**

<sup>1</sup> See e.g. Craig (1990).

<sup>2</sup> See, especially, Timothy Williamson (2000), for discussion of the claim that knowledge is the norm of assertion, and John Hawthorne (2004) for the claim that knowledge is sufficient for use in practical reasoning.

<sup>3</sup> See Lewis (1980: 263–93). He later abandons the Principle Principal in Lewis (1994: 473–90).

<sup>4</sup> See Lewis (1996: 549–67).

<sup>5</sup> I am thinking of the particularist approach of Jonathan Dancy (2004).

<sup>6</sup> The initial account of the fallacy can be found in Robert Shope (1978: 397–413). Some of the lessons of the fallacy for various counterfactual accounts in ethics are detailed in Robert N. Johnson's excellent piece ‘Virtue and Right’ (2003: 810–34).

<sup>7</sup> See e.g. Michael Smith (1995).

<sup>8</sup> See ch. 5 of Kvanvig (2003). This conclusion involves a change from the conclusion reached in Kvanvig (1998: 426–51), where I claimed that the Swamping Problem undermined every version of externalism as well as undermining nearly every version of internalism (except for a radically subjective one).

<sup>9</sup> Kvanvig (2003), see especially ch. 1.

<sup>10</sup> It should be noted that Pritchard does not hold that an appeal to final value can solve the value problem. The passage quoted reports an approach to the special and unique value of knowledge that Pritchard investigates but in the end rejects.

<sup>11</sup> Thanks to Christian Piller for helping me to see this point clearly.

<sup>12</sup> The literature includes accounts that explicitly claim to be such, as well as the buck-passing accounts that have become more popular since the publication of T. M. Scanlon's (1998) *What We Owe to Each Other*. For the response-dependent approaches, see e.g. Mark Johnston (1989), David Lewis (1989), Michael Smith (1995), John McDowell (1985), David McNaughton (1988), and David Wiggins (1987).

<sup>13</sup> Robert Shope (1978: 397–413).

## **6 From Epistemic Expressivism to Epistemic Inferentialism**

<sup>1</sup> See for instance Unger (1984), Cohen (1988), DeRose (1992, 1995, 2009), Lewis (1996), Rysiew (2001), Neta (2003), Hawthorne (2004), Stanley (2005).

<sup>2</sup> See for instance the virtue epistemology of Sosa (1980, 2007) and Zagzebski (1996), but also those who think knowledge is a norm of belief, action, and/or assertion, e.g. Pollock and Cruz (1999: ch. 5), Williamson (2005), Hawthorne and Stanley (2008), and those who follow Sellars (1956) in thinking that knowledge attributions 'place one in the space of reasons', e.g. Williams (1992, 2001), Brandom (2000a: ch. 2), and Rosenberg (2002).

<sup>3</sup> In Bar-On and Chrisman (2009) and Chrisman (2009: §5), non-standard forms of expressivism are explored, which attempt to be semantically neutral. However, in this paper about the meaning of normative claims, such as epistemic claims, I set those views aside as not addressed to the present issue.

<sup>4</sup> See Grice (1957), Schiffer (1972), Davis (2003).

<sup>5</sup> Compare Schroeder (2008) and Chrisman (forthcoming) for more discussion of the semantic commitments of expressivism. Early expressivists like Ayer and Carnap can be plausibly read as denying that ethical terms have semantic value at all. However, this has come to seem increasingly implausible, which is why most recent expressivists should be read, I believe, as relying on an ideationalism as their general semantic approach.

<sup>6</sup> His is, at its core, a theory of judgements about *rationality* rather than a theory of judgements about something being *good*. I'm simplifying for present purposes. What he actually writes is this: 'An observer, we may assume, accepts at least an incomplete system of general norms; call that system *N*. He then thinks an alternative rational if and only if he thinks it *N-permitted*—that is to say, rational according to *N* . . . whether an alternative is *N*-permitted is a matter of fact, and so a belief that the alternative is *N*-permitted is a factual belief. Thinking X rational, then, is a combination of a normative state and a state of factual belief' (Gibbard 1990: 91).

<sup>7</sup> See especially Davidson (1980: ch. 5), Smith (1987, 1994: ch. 4).

<sup>8</sup> See Mackie (1977: ch. 1).

<sup>9</sup> Some expressivists such as Ayer (1946), Stevenson (1937), and Gibbard (1990) have assumed that this commits them to denying the truth-aptness of ethical claims. However, more recent expressivists such as Blackburn (1998), Timmons (1999), Gibbard (2003), and Ridge (2009) have sought to avoid this commitment in various ways.

<sup>10</sup> One exception to this is ethical expressivists pursuing the 'quasi-realist' programme of reconstructing the language of realism within an anti-realist expressivism. Thus, both Blackburn (1996: 87, 1998: 318) and Gibbard (2003: 235) argue briefly for extending their ethical expressivism to cover epistemic claims. Outside of that, Field (1998) and Heller (1999) have sketched views of particular epistemic claims that might be read as expressivist views, although they have not been worked out in very much detail.

<sup>11</sup> Compare citations in note 2 above.

<sup>12</sup> See for instance Kvanvig (2003), Pritchard (2007b), and the papers collected in Haddock, Millar, and Pritchard (2009).

<sup>13</sup> Kvanvig (2003: ch. 7), Cuneo (2008: chs 5–6), and Lynch (2009a) have argued that epistemic expressivism faces special problems not faced by ethical expressivism. These arguments are discussed and rejected in Carter and Chrisman (unpublished).

<sup>14</sup> For example, Copp (2001) defends a view he calls ‘realist expressivism’, where the expressive function of ethical claims is seen as deriving from something other than their semantic content. See also the papers cited in note 3 above.

<sup>15</sup> In addition to the quote from Gibbard in note 6 above, compare Gibbard (1990: 93) and Ridge (2006) for the most explicit examples of this strategy. Both, however, echo Hare's (1952: 18–20) distinction between the neustic and phrastic elements of meaning. One possible exception to this is Blackburn (1988, 1993), who argues that logically complex sentences serve a particular commissive function; that is, they ‘tie one to a tree’ of conditional commitments. For example, the claim ‘If your stealing that money was wrong, then you will eventually be punished’ commits its author to the content of its consequent if she is committed to the content of its antecedent, and to the negation of its antecedent if she is committed to the negation of its consequent. I think approaches of this sort are quite promising; however, they seem to me to rest on a tacit abandonment of the ideationalist alternative to truth-conditional semantics, from which expressivism seemed to receive some shelter from objections to its implicit semantic assumptions. For to hold that logically complex statements involving normative predicates mean what they do in virtue of ‘tying one to a tree’ of conditional commitments is to endorse something much more like a conceptual or inferential role semantics. If that's right, then the resulting view may be consistent with the advantages traditionally claimed for expressivism (more on this below), but I think it can no longer be fairly thought of as a species of the view that ethical claims mean what they do in virtue of the mental states that they express.

<sup>16</sup> One possible exception is the version of expressivism explored in Schroeder (2008), which seeks to reconstruct the compositionality and semantic properties of normative statements in a general expressivist framework. However, in order to capture the semantic properties, he ends up forcing the expressivist to reconstruct notions of validity and entailment for all statements (both normative and non-normative) in a global expressivist account of language—that is, all indicative sentences are thought to mean what they do in virtue of the desire-like attitude of ‘being-for’ that they express. Even if this view escapes the criticism I go on to levy in the text above, it is not a very plausible view. As Schroeder himself concludes, it is ‘an extremely unpromising hypothesis about the workings of natural language’ (2008: 179).

<sup>17</sup> This is a simplification of a distinction deployed in Brandom (2009) drawing on Sellars (1956).

<sup>18</sup> For philosophical expositions of theories of concepts which are at least inspired by this general tradition, see Dretske (1981) and Fodor (1990), though both Dretske and Fodor pursue more complex models, anticipating many of the most natural objections to the crude summation of the tradition presented in the text above.

<sup>19</sup> For philosophical expositions of theories of concepts which are inspired by this general tradition, see Sellars (1956), Dummett (1974), Rosenberg (1974), and Brandom (1994, 2000a, 2009).

<sup>20</sup> In this way, the account bears some similarity to Wedgwood's (2007: chs 4–5) conceptual-role semantics for normative claims. However, I differ from Wedgwood in thinking that conceiving of meaning in terms of inferential-roles opens up room for anti-realist views in normative ontology rather than forces one to adopt a very strong form of normative realism.

<sup>21</sup> I leave it open here whether, like truth-conditionalism, inferentialism should be seen as a global view of meaning. If it is, normative inferentialism is an application, which draws a contrast between the normative and the non-normative in terms of type of inferential role.

<sup>22</sup> You might also be able to answer in the order of causes by saying 'Because I approve of giving to charity'. This exhibits the ambiguity of the question 'Why did you do X?' Ethical inferentialism as I am conceiving of it is consistent with this response in the order of causes, but it denies that such a response is adequate in the order of reasons.

<sup>23</sup> At least, this could be viewed as a necessary condition for such commitment. Things may be more complicated depending on one's view of ontological commitment.

<sup>24</sup> What about views which see ethical claims as having both practical and theoretical inferential connections? On the present account, as long as a claim like 'X is good' commits one to 'I shall promote X', it has a practical inferential role and so does not commit one ontologically to the fact that X is good. However, it might imply that X is approved of by many people, and this then could commit one to the fact that X is approved of by many people.

<sup>25</sup> Compare Kelly (2003) and my discussion of Kornblith (2001) in Chrisman (2008: 354–8).

<sup>26</sup> Of course, on some accounts of assertion, asserting that *p* commits one to 'I know that *p*'; however, even if that's right, it's still unclear that the implicit knowledge claim is providing the link to motivation rather than the explicit claim that *p*.

<sup>27</sup> Consider the debate over whether knowledge is a ‘norm of action’ among, for instance, Pollock and Cruz (1999: ch. 5), Williamson (2005), Stanley (2005), Hawthorne and Stanley (2008), Neta (2008), and Brown (2008).

<sup>28</sup> These ideas are related to Austin's suggestion that ‘When I say “I know”, I give others my word’ (1946, 1979: 99). However, I take it Rorty and Craig want their ideas to extend to third-personal ascriptions of knowledge as well.

<sup>29</sup> I would like to thank an audience at the University of St Andrews as well as J. Adam Carter, Georgi Gardiner, Graham Hubbs, Alan Millar, and Michael Ridge for helpful feedback on earlier versions of this material.

## **7 Norms of Trust**

<sup>1</sup> Elster (1989: 100).

<sup>2</sup> A point observed by both Austin (1946) and Anscombe (1979).

<sup>3</sup> Adler labels this norm the default rule: ‘one ought simply to accept a speaker's testimony unless one has special reason against doing so’ (2002: 143). And he observes that the ‘default rule actually functions as a *presumption* that our informants are being cooperative’ (2002: 154).

<sup>4</sup> This point is observed by Adler: ‘Testimonial situations allow modelling as repeated Prisoner's Dilemmas at least, presumably, in their origins (when testimony is more functional and local, and so defection more readily detected)’ (2002: 155). And Pettit claims that the norm of ‘Telling the truth reliably rather than expediently, randomly, or whatever’ is ‘equivalent to cooperating in a many-party prisoner's dilemma’ (1990: 735).

<sup>5</sup> See Faulkner (unpublished).

<sup>6</sup> It is developed in detail in Axelrod (1984). See also Blais (1987).

<sup>7</sup> I follow the terminology of Sripada and Stich (2006).

<sup>8</sup> I discuss Williams's genealogy in more detail in Faulkner (2007b).

<sup>9</sup> I tried to work out a definition of this thick sense of trust in Faulkner (2007a).

<sup>10</sup> This explication is closer to the surface when more is at stake. Maybe it is a first date and *A* and *S* are enamoured, or *A* is a potential informant and *S* is trying to gain his confidence in a political climate where both have much to lose.

<sup>11</sup> See Adler (2002: ch. 5).

<sup>12</sup> This chapter was written whilst in receipt of AHRC funded research leave. My thanks to Adrian Haddock and an anonymous referee.

## **8 Testimonial Entitlement and the Function of Comprehension**

<sup>1</sup> Positive epistemic statuses make up the subject matter of epistemology. Some apply to beliefs, some to believers. Knowledge and truth are both positive epistemic statuses that apply to belief. Warrant is another. The connections between knowledge, true belief, entitlement, and justification are complex. Justifications and entitlements both fall short of enabling knowledge, even when they warrant a true belief. Such is the stuff of Gettier cases. Justification (as understood here) is clearly not necessary for perceptual knowledge (Kornblith 2008). I can only explicate testimonial entitlement here, and cannot pursue its connections with testimonial justification and testimonial knowledge. I have discussed testimonial knowledge elsewhere (Graham 2000a, 2000b, 2006b). See note 4 for my construal of the adjective ‘testimonial’.

<sup>2</sup> ‘*Prima facie*’ entitlement means defeasible entitlement. Counter-considerations may overturn the entitlement. Justification or further entitlements may be required to restore the warranting force of an entitlement. ‘*Pro tanto*’ entitlement means *some* entitlement, where some may fall short of enough for on balance entitlement, given some threshold (Graham 2006a). Justification or further entitlement may supplement the entitlement, converting some warrant into enough. Do not confuse my thesis that comprehension-based beliefs generally enjoy *prima facie pro tanto* entitlement with the obviously absurd thesis that comprehension-based beliefs *ipso facto* enjoy indefeasible on balance warrant, that believing everything you hear somehow confers maximal epistemic warrant.

<sup>3</sup> I say ‘a kind of entitlement’ because I believe one kind derives from normal functioning (as I argue below) and another related kind may derive from function fulfilment. Still a third related kind may derive from function fulfilment due to normal functioning. This third kind is related to what Duncan Pritchard would call an achievement, a success through the exercise of ability. On a looser understanding of ‘achievement’ where

achievements are just successes or fulfilments, both the entitlement that turns on normal functioning and the entitlement that turns on function fulfilment would be achievements. I hope to pursue the connection between kinds of entitlement and notions of achievement elsewhere.

Michael Bergmann (2006) (following Plantinga 1993), Tyler Burge (2003), and Ernest Sosa (2007) each develop accounts of entitlement structurally similar to mine. My account explicitly relies on the etiological notion of function and normal functioning. Plantinga and Sosa explicitly reject reliance on etiological functions. Burge remains neutral on etiological functions and on the importance of normal functioning. I intend to discuss their views in detail elsewhere.

<sup>4</sup> My subject matter is the epistemology of beliefs based on our capacity to comprehend assertive speech acts. This category is much broader than the category of testimony-based beliefs, insofar as (a) we take the latter category to involve a narrow account of testimony (as a speech act), and (b) we require that a testimony-based belief that *P* be caused or sustained by a speaker's testimony that *P*. Many comprehension-based beliefs do have, as their distal causes, speakers having testified that such and such is the case. But many do not. For testifying (taken narrowly) is just one of many ways a speaker can assertively present a proposition (Alston 2000; Bach and Harnish 1979); testifying is a particular way of assertively presenting a proposition, involving additional intentions on the part of the speaker (Graham 1997). So there are plenty of comprehension-based beliefs that are not 'based' on a speaker's testimony. Furthermore, beliefs based on miscomprehending a speech act are still comprehension-based. A speaker may question whether *P*, but the hearer miscomprehends the force, and takes the question to be an assertion, and comes to believe that *P*. Or a speaker may assert *Q*, but the hearer miscomprehends the content, and comes to believe that *P*. Since both beliefs are exercises of the hearer's general capacity to comprehend speech, they are both comprehension-based beliefs. The category of comprehension-based beliefs is thus larger than the category of testimony-based beliefs, and larger than the category of assertion-based beliefs. The epistemology of 'testimony' literature is, for the most part, interested in this broader category. 'Testimonial entitlement' is thus entitlement for beliefs based on comprehension.

<sup>5</sup> Representing an utterance as the assertion that *P* sometimes leads a hearer to form the belief that the speaker believes *P*, or to believe that the speaker asserted that *P*. But the central case of a comprehension-based belief is a belief in the very propositional content taken as asserted. For discussion of this point, see Hunter (1998).



I allow for leeway as to what counts as the propositional content taken as asserted. It may include simple logical entailments or obvious evidential consequences from what was strictly speaking said or meant. If someone says John was at the party, you may take the speaker to have also asserted that someone was there (Goldberg 2001). It may include obvious conversational implicatures. Someone says there is a gas station around the corner, and you take the speaker to assert that you can fill your tank there. It may also include what was communicated non-verbally, given an already existing linguistic context. Pointing in response to a query is a perfectly good way of telling someone where something is, or where to go. Assertions are also made both directly and indirectly. You can assert by asking a question or making a command. In such a case, you make two speech acts with one utterance.

<sup>6</sup> For a survey of impressive psychological evidence in favour of the claim that credulity is the default, see Gilbert (1991). For additional evidence and argumentation, see Gilbert (1993). Gilbert examines three models of the transition from comprehension to belief: Cartesian, Spinozan, and Cartozan. The Cartesian sees comprehension as neutral towards belief. The Spinozan sees comprehension as consisting in belief; to comprehend an assertion as an assertion that *P* just is to believe *P* as a result of one's linguistic ability to comprehend. The Cartozan sees comprehension states as disposing belief, but still ontologically distinct from belief. He favours the Spinozan view. The research he surveys, and his own research conducted with colleagues, clearly undermines the Cartesian view. But it is not obvious that it favours the Spinozan over the Cartozan position. I favour the less controversial, Cartozan view. I discussed Gilbert's view in Graham (1999).

<sup>7</sup> Witness Origgi and Sperber: 'There are . . . reasons to expect the ability to attribute false beliefs to others to develop after the ability to communicate verbally. The attribution of false beliefs to others plays an obvious role in the ability to filter false information communicated either by mistaken or deceitful speakers' (2000: 163)

<sup>8</sup> There is the use of 'function' to describe how something works or operates—what it does—as a part of a larger system (Cummins 1975). Functions, in this sense, are the causal role capacities of parts that contribute to some capacity of the containing system. If the part lacks the capacity, it lacks the function. A 'functional analysis' of the system reveals 'functions'—capacities or dispositions—of the parts; the analysis reveals how the system operates or 'functions'. Call these causal-role functions, or just c-functions.

<sup>9</sup> The etiological account was inspired by Larry Wright's (1973) article 'Functions' and then developed independently along different lines by Millikan (1984, 2002), Karen

Neander (1991a, 1991b), Peter Godfrey-Smith (1994), and David Buller (1998, 2002), among others. For a useful anthology, including challenging criticisms of the etiological view, see Buller (1999).

<sup>10</sup> There are (at least) four attractive features of the etiological account of functions. First, as we have just seen, it gives an informative and non-circular account of normal conditions and normal functions. Second, it underwrites teleological explanations; something is there because of its function. Third, it accounts for the difference between the function of an item and side effects that may be beneficial for the item or the user of the item. The function of my nose involves affecting incoming air, but not to hold up my glasses. The function of a large belt buckle in the Old West is to keep up trousers, and maybe also to display social status, but not to stop bullets during a shoot-out at the corral. Fourth, it explains malfunction. A diseased or deformed heart still has the function of pumping blood even if it cannot, for the heart is there because its ancestors pumped blood. A broken belt buckle is still supposed to hold up trousers, for it was reproduced on the model of buckles that did.

Another account of functions aimed at accounting for biological functions is the propensity account (Bigelow and Pargetter 1987). On this account the function of an item is its propensity to contribute to the fitness of the organism in its natural habitat. There are four problems with this account. First, the view helps itself to an unexplained account of natural habit. Second, it fails to support teleological explanations, for propensities are future effects, and future effects cannot explain the current existence of anything. Third, it does not capture the function/side-effect distinction: holding up glasses may have a propensity to contribute to fitness, but that is not a function of my nose. Fourth, malformed hearts have no propensity to contribute to fitness, so they have no function on this view (cf. Neander 1991a, 1991b).

<sup>11</sup> Since an item may fail to fulfil its function even while in normal conditions while functioning normally, it is a mistake (too easily made) to suppose that failure to fulfil a function implies that the item is either not in normal conditions or not functioning normally (malfunctioning).

<sup>12</sup> See, for example, Graham (unpublished).

<sup>13</sup> E.g. McDowell (1980), Millikan (1984), Sperber (2001), Origgi and Sperber (2000), Faulkner (2000), and Schmitt (2006).

<sup>14</sup> This argument appears throughout Millikan's writings (cf. 1984: 30–2, 52–60, 2004: 25–6, 2005: 14, 16, 44–5, 57–8, 90, 94–5, 153, 157, 169, 172).

<sup>15</sup> Though Dan Sperber (rightly) doubts assertive communication will be as reliable as perception, he agrees with Millikan on evolutionary grounds: ‘For communication to stabilize within a species, as it has among humans, both the production and the reception of messages should be advantageous. If communication were on the whole beneficial to producers of messages (by contributing to their fitness) at the expense of receivers, or beneficial to receivers at the expense of producers, one of the two behaviors would be likely to have been selected out, and the other behavior would have collapsed by the same token (incidentally, there are exceptions, in particular in inter-species communication). In other words, for communication to evolve, it must be a positive-sum game where, in the long run at least, both communicators and receivers stand to gain’ (Sperber 2001: 403). For further agreement and discussion of the stability of truth-telling, see Sober (1994a).

<sup>16</sup> A casual view of hearer recognition of speaker's communicative intentions might suppose that it consists in the (true) conscious explicit belief that the speaker intends the hearer to recognize his (the speaker's) informative intention. I seriously doubt whether the recognition consists in such a belief. One reason for this is that young children comprehend assertions—and thereby recognition communicative intentions (if the Gricean account is correct about communication)—but they do not have the capacity to form such beliefs. See also Hunter (1998).

<sup>17</sup> The *argument from occasion specific intentions* and the *argument from communicative success* are clearly related. The first says the function of a particular assertion is whatever intention the speaker has in mind in making the assertion. The second says the function of communication is the recognition of a particular intention on the part of the speaker, an intention that is always present when an assertion is made, always recognized when communication occurs. The function of assertion on both arguments is just the fulfilment of a speaker's intention. As Millikan has noted (2005: 28), tying function to intention is a dominant, traditional philosophical view about how communicative functions are determined. This view has wide appeal outside of philosophy as well. Witness Michael Tomasello: ‘The notions of communicative intention and function are correlative. Someone uses a piece of language with a certain communicative intention, and so we may say that that piece of language has a certain function’ (2003: 3). But since we know that functions from histories are one thing and functions from intentions are another, we know that neither of these arguments touches the thesis that the etiological function of assertion is to induce true hearer belief. For functions from intentions, look to intentions. For functions from histories, look to histories.

<sup>18</sup> John McDowell has argued that conveying knowledge is the function of communication (1980). Though obviously congenial to my point of view, McDowell's

paper lacks (a) an explication of functions, especially etiological functions and how they differ from functions from intentions, (b) a discussion of normal functioning, (c) the distinction between a one-off effect as function, and reliability vis-à-vis that effect as function, (d) an extended discussion of why conveying knowledge should be the function, and (e) and explanation of the relevance of the function of communication to the epistemology of beliefs based on comprehension. McDowell also focuses on knowledge from comprehension, and does not isolate entitlement as a distinctive epistemic achievement.

<sup>19</sup> Speaking of cooperative communication in general, and not just human linguistic communication, Marc Hauser writes: ‘The last 20 years or so in biology has taught us . . . that whenever individuals assume different roles in an apparently cooperative venture, there are opportunities for individuals to cheat, reaping the benefits without paying the costs’ (2000: 141). It’s a general strategy. He continues: ‘some animals have the capacity to detect mismatches between a signal’s common function and the way it is currently being used. These animals can therefore detect cheaters, liars, and other self-interested deceivers. They have the mental tools to see through deceit’ (2000: 156). We are animals like these.

<sup>20</sup> This chapter was long in the writing, from conception to completion. I am grateful to Duncan Pritchard for his invitation to participate in the conference where some of the ideas in the chapter were presented in their infancy. At the conference I benefited from Finn Spicer’s enthusiastic commentary and from questions from Sandy Goldberg, Alvin Goldman, Jon Kvanvig, and Ernest Sosa. A later presentation to a workshop at UCLA helped move it along. Thanks to Sven Bernecker, Tim Black, Peter Kung, Nikolaj Pederson, and Masahiro Yamada for their questions on that occasion. Discussions with Justin Coates and Patrick Todd spurred completion. The comments from a reader for Oxford University Press led to some last-minute improvements.

## **9 Knowing from Being Told**

<sup>1</sup> Angus Ross (1986) and Richard Moran (2006) both resist the idea that acts of telling should routinely be viewed as constituting evidence for the thing told for reasons associated with the fact that acts of telling are acts freely undertaken by the speaker. My conception of evidence is more minimal than theirs and I agree with them that the

epistemology of testimony should not be assimilated to the epistemology of empirical evidence in general.

<sup>2</sup> Such a view is treated sympathetically in Fumerton (2006). See especially p. 80.

<sup>3</sup> The epistemological framework for the present discussion is developed at greater length in my contribution to Pritchard, Millar, and Haddock (2010). There is some overlap between this chapter and the last of the chapters in that contribution, though the treatment of testimony here is somewhat fuller.

<sup>4</sup> Acts of saying that *p* should be taken to include cases in which the informant simply says, 'Yes', when asked whether *p*. Likewise, saying, 'No', in response to an enquiry whether *p*, should be taken to be an act of saying that not-*p*.

<sup>5</sup> In what follows I draw on an account of practices developed in Millar (2004).

<sup>7</sup> When, as in Williamson (2000), it is said that there is a knowledge-rule for assertion—assert that *p* only if you know that *p*—is assertion being understood broadly to encompass declarative utterances that include advice and expressions of opinion or more narrowly so that it is close to what I am calling telling? If the former then it seems to me that there is no knowledge-rule for assertion. But what I am calling telling might be what some theorists call assertion.

<sup>8</sup> An example of this sort is considered in Welbourne (2001: 102).

<sup>9</sup> Further refinements are required. Practices do not exist in isolation and so wider considerations might dictate that one sticks with a practice while flouting the rules. (See Millar (2004: ch. 3)). My thinking on commitments is much influenced by John Broome's work. See, for instance, Broome (2002). However, his larger framework should not be read into mine and I do not suggest that he would endorse the ideas presented here.

<sup>10</sup> Compare McDowell (1984: esp. sections 11–12). Mulhall (1990: especially ch. 4) is instructive on the problems of explaining the understanding of utterances, on the model of interpreting sounds conceived as devoid of meaning.

<sup>11</sup> When corrected for misusing the term 'arthritis', Burge's (1979) patient does not respond by saying that what he means by 'arthritis' is different. He adjusts his usage.

<sup>12</sup> I put the notion of recognitional ability to work in Millar (2007a, 2007b, 2008), and in my contribution to Pritchard, Millar, and Haddock (2010).

<sup>13</sup> But see M. Fricker (2003).

<sup>14</sup> Richard Fumerton thinks that the following is ‘probably right’: ‘The meanings of sounds or marks are to be found in the head, or better in the minds, of those who produce those sounds and marks. To interpret reasonably those sounds or marks we now need to solve the problem of other minds. . . . When the language we are interpreting is our own—without begging the question, when the sounds or marks appear to be of the same type as those we ourselves use to express thoughts—one might suppose that our best hope is to rely on some version of an argument from analogy or an argument to the best explanation’ (2006: 83).

<sup>15</sup> On this see McDowell (1981, 1982). For related discussion of the scope of perceptual knowledge, see Millar (2000).

<sup>16</sup> The importance of discrimination for perceptual knowledge is emphasized in Goldman (1976). Here I am emphasizing its importance for knowledge of what phenomena indicate.

<sup>17</sup> This is in keeping with Miranda Fricker's (2003) emphasis on the importance of a testimonial sensibility.

<sup>18</sup> This is the truth behind, what might be called, *defaultist* accounts of the epistemology of testimony. J. L. Austin, for instance, writes, ‘It is fundamental in talking (as in other matters) that we are entitled to trust others, except insofar there is some concrete reason to distrust them’ (1946, 1979: 82). H. H. Price discusses the maxim: ‘Believe what you are told by others unless or until you have reasons for doubting it’ (Price 1969: 124). Tyler Burge advances the principle that we are ‘entitled to accept as true something that is presented as true and that is intelligible to [us], unless there are stronger reasons not to do so’ (1993: 467). Michael Welbourne says that ‘the [by implication appropriate] *default* response to testimony is to accept what we are told’ (2001: 104–5). For criticism of this type of position, see E. Fricker (1994). For criticism of Fricker, see Goldberg and Henderson (2006). For Fricker's defence, see Fricker (2006). My own position does not invoke the idea that the appropriate default stance to testimony is acceptance. I doubt that an adequate epistemology of testimony is delivered solely by highly general considerations about speech acts or about presuppositions of rational communication.

<sup>19</sup> Elizabeth Fricker (1994) thinks of trustworthiness as having these two components. Much of what I say is in the spirit of her account though, unlike her, I am not working within the framework of a conception of knowledge as built out of justified true belief.

<sup>20</sup> This is more fully explored in Millar (2008).

<sup>21</sup> I borrow the example from Travis (2005), which has helped to shape my own thinking.

<sup>22</sup> This sort of case is discussed by Jennifer Lackey (2007) who takes it to yield knowledge.

<sup>23</sup> Jonathan Adler emphasizes that ‘[o]nce testimony is accepted, there are many sources which come into play which supply additional grounds relevant to the warrant for the belief’ (1994: 266). That is surely right. My point is rather different since it concerns the role of repeated encounters with sources of information, not in furnishing us with more evidence supporting something acquired from testimony, but in inculcating factual knowledge that amounts to little more than an ability to recall or acknowledge a publicly available known fact.

<sup>24</sup> The kind of knowledge under discussion is more fully explored in my contribution to Pritchard, Millar, and Haddock (2010) under the heading ‘detached standing knowledge’.

<sup>25</sup> I am grateful to the audiences at the Stirling Social Epistemology conference and at the University of Manchester for helpful discussion of the issues addressed in this chapter. I have benefited enormously from frequent discussions with Adrian Haddock, Duncan Pritchard, and colleagues at Stirling. Thanks are due to the University of Stirling for research leave during which this chapter was initially drafted and to Department of Philosophy for further support.

## **10 Can A Priori Entitlement be Preserved by Testimony?**

<sup>1</sup> See Burge (1993, 1997, 1999).

<sup>2</sup> See e.g. Christensen and Kornblith (1997) and Malmgren (2006). The anti-Burgean view that Christensen and Kornblith and Malmgren all defend—the view that all testimonially preserved warrant is a posteriori—is a widely received view in the epistemology of testimony.

<sup>3</sup> For Burge's articulation of this teleological picture, see Burge (2003: 505–6). Burge acknowledges the great similarity between his own teleological picture and those offered by Sosa (1991) and Plantinga (1993).

<sup>4</sup> ‘I think that any being that has entitlements to beliefs must have justifications for some other beliefs—or at least must be capable of having justificatory support for some other beliefs. For I think that having beliefs requires being able to carry out inferences, and the

relevant inferences must sometimes support beliefs. Inferentially supported beliefs are justified, or are at least backed by beliefs that are capable of justifying them. Higher non-linguistic animals have reasons or justifications for some of their beliefs in this sense' (Burge 2003: 504–5 n.1).

<sup>5</sup> See the last sentence of the passage from Burge (2003) quoted in note 4 above.

<sup>6</sup> See especially his contrast between the preservative and substantive role of memory in Burge (1993: 462–5).

<sup>7</sup> 'I use the term "presentation as true" to cover more than assertions and judgments. Obvious presuppositions, or conventional implicatures, are examples. When someone says to kill the shortest spy, he or she presents it as true that there is a shortest spy. In such cases, as well as the indicative cases, the entitlement to accept what is presented as true can be independent for its justificational force of perceptual connection to context' (Burge 1993: 482–3).

<sup>8</sup> Of course, deaf people may communicate by means of visual sign language. And perhaps it is possible for creatures to communicate by means of some extra-sensory perception. The only feature of testimony that matters, for Burge's view, is that there is a 'passage of propositional content from one mind to another' (Burge 1993: 481), however that may happen.

<sup>9</sup> Since warrants (both entitlements and justifications) are typically, for Burge, defeasible, I will drop the phrase 'prima facie' when speaking of warrants throughout this paper: warrants are all to be understood as prima facie, unless otherwise specified.

<sup>10</sup> 'It is understanding whose exercise in particular instances does not require in those instances perceptual warrant for the application of what is understood. A first approximation elaboration is that it is conceptual understanding that does not require, in thinking and understanding an intentional content, perceptual warrant for the *de re* application of some aspect of the content' (Burge 1997: 21–2).

<sup>11</sup> Burge (1997: 22).

<sup>12</sup> I am grateful to Yuval Avnur and Alan Millar for helpful comments on an earlier draft of this chapter.



## 11 The Assurance View of Testimony

<sup>1</sup> Epistemic justification is whatever kind is required for knowledge in those instances of knowledge for which it requires justification.

<sup>2</sup> For our purposes in this chapter, it does not matter just what testimony or testimonial justification is. I address the question whether assurance is necessary or sufficient for giving a non-evidential justification in cases we commonly call testimonial. To this end, we can use the broadest disjunctive sense of ‘testimony’. For a good presentation of a broad, ‘disjunctive’ view of testimony, see Lackey (2008: ch. 1).

<sup>3</sup> A non-reductive view would take testimonial epistemic reason-giving or justification to be in part *sui generis*. A non-individualist view would take testimonial reason-giving or justification to depend in part on conditions other than the belief *p*'s being based on the epistemically justified beliefs of the hearer. Examples of non-individualist views are the transindividual reliability view, that testimonial epistemic justification depends a reliable process of testimonial propagation and hearer uptake (Schmitt 1999; Goldberg 2007a: esp. 200–33), the social responsibility view, that it depends on the testifier's being epistemically responsible in believing *p* (Owens 2000, 2006), and the transindividual reasons view, that it depends on the testifier's possession of an epistemically good reason to believe *p* (Schmitt 2006).

<sup>4</sup> The view traces back to Ross (1986). Versions have been developed by Hinchman (2005, forthcoming) and Moran (2005, 2006, unpublished). I focus here on Moran's (2006). All page references to Moran not otherwise marked are to that article.

<sup>5</sup> I take it that accepting the testifier's assurance that *p* is not the same as *trusting* the testifier that *p*. Accepting is an act and entails an acknowledgement of the assurance and a recognition that one accepts. Trusting is a disposition or a practice and does not entail any acknowledgement of the assurance or even a recognition that one trusts. An assurance view that requires the addressee's uptake for reason-giving should maintain that accepting, not trusting, is necessary for reason-giving. Trusting may not always be voluntary, but accepting is on an agreement version of the assurance view. This raises a concern about this view—that in some instances in which assurance gives a reason, the addressee cannot be said to accept because uptake is involuntary. However, I will not press this concern.

<sup>6</sup> I use ‘propositional’ and ‘doxastic’ in application to justification in the senses familiar from the literature stemming from Firth (1978).

<sup>7</sup> I take it that assurance or its acceptance counts as giving an *epistemic* reason to believe *p* just in case it gives a reason that, were it possessed by the addressee and good, would

contribute to the prima facie epistemic justification of the addressee's belief  $p$ . As far as I can make out, Moran's formulation of the assurance view presupposes that assurance or its acceptance counts as giving an epistemic reason to believe  $p$  just in case it gives a reason that is *epistemically assessable by the addressee*, where this epistemic assessment involves judging that the reason has an epistemically relevant feature. (That Moran takes assurance to give a reason only if the reason it gives is epistemically assessable by the *addressee*, as opposed to a third party evaluator, is shown by his remark that 'the speaker's knowledge and trustworthiness is epistemically inert for the audience until the question of the particular speech act or illocution is settled' (Moran 2006: 289). The word 'settled' here can only mean 'answered by the audience'.) For now I will treat the account of giving an epistemic reason expressed in the first sentence of this note and Moran's presupposition as equivalent, but where they come apart, the former must be preferred if we are to relate the assurance view to the issue of non-reductive testimonial epistemic justification that has driven the testimony literature.

<sup>8</sup> A remark of Moran's (2006: 272–3) suggests that he prefers the assurance view to the reliability view of testimonial justification on the ground that testimony is often unreliable. But this makes no sense. His own assurance view yields the non-sceptical outcome that we have justified testimonially based beliefs only when the background conditions of trustworthiness and reliability are satisfied. He cannot consistently support his view as a non-sceptical alternative to the reliability view by appeal to the frequent unreliability of testimony.

<sup>9</sup> Moran identifies assuring that  $p$  with telling that  $p$ . But the relation between the two is more complicated than he notes, since assurance often (if not always) comes in degrees and telling never does. Possibly, telling is assuring to a sufficient degree. Are we to take it that any degree of assurance involves the conditions Moran proposes for telling: a presentation of the utterance as giving a reason, an offer of responsibility, and the like? Possibly any sufficient degree of assurance does.

<sup>10</sup> Hinchman (2005) differs from Moran in claiming that telling requires the addressee's recognition that the teller, in asserting that  $p$ , intends the addressee 'to gain access to a prima facie entitlement to believe  $p$  through this recognition' (2005: 567). Moran does not expressly deny that telling requires an intention to give the addressee a reason to believe, nor that giving the addressee a reason to believe depends on a recognition of this intention. But he does argue persuasively (Moran 2006: 284–92) that the recognition of such an intention is not *sufficient* to give a reason. It seems to me that an assurance view should also reject Hinchman's claim that an intention or its recognition is *necessary* for conferring access to an entitlement to believe: it should be enough that by inviting the

addressee to trust or offering to take responsibility, the addressee gains access to an entitlement. An interpersonal offer, rather than a personal intention, should be sufficient for making a reason available. The intention is at most a condition of the reason's being *good*. If an intention is necessary for telling, then telling is not necessary for the sort of assurance that suffices for making a reason available.

<sup>11</sup> See note 3 for formulations of these views.

<sup>12</sup> That Moran endorses this clause is indicated, for example, by the paragraph where he asks how the 'speaker's words . . . acquire the status of *a reason to believe something*' (288, emphasis deleted and added).

<sup>13</sup> That Moran endorses this clause is indicated, for example, by the paragraph where he says: 'the idea is not that the speaker's words 'all by themselves' should count as a reason for belief . . . As with the explicit assumption of responsibility that goes with making a promise, its success will depend on the various conditions that go into the speaker's being in a position to take on any such responsibility' (289). Though Moran speaks at first of 'a reason for belief', the 'success' to which he refers in the second sentence is most naturally taken as success in giving a *good* epistemic reason. In this case, he is saying that giving a good epistemic reason depends on the speaker's being in a position to take on a responsibility. This fits the analogy with promising: a promise gives a practical reason to rely on the promised action, but the promisor must be in a position to carry through or take the consequences of not doing so if the practical reason for reliance is to be a good one.

<sup>14</sup> Moran speaks of 'responsibilities with respect to the belief in question'. I assume that he really intends that the testifier is epistemically responsible in *asserting p*, rather than in believing *p*. For the testifier need not believe *p* for the testimony to give the addressee a reason to believe *p*.

<sup>15</sup> One might suggest that for Moran the assurance or its acceptance does not give the epistemic reason but merely determines that the reason given (in virtue of the obtaining of the background conditions) is epistemic. Some of his remarks suggest this alternative view. I am not sure how to understand the difference between this view and the one I attribute to Moran in the main body of the text. It seems that the objections I raise to the view I attribute would apply to this view as well, if it really does differ from the former view.

<sup>16</sup> Candidates for such norms include the Norm of Veracity, that one has a prima facie obligation to assure that *p* only if one takes *p* to be true, and the Norm of Credulity, that it is prima facie permissible to trust any proposition one is assured of.

<sup>17</sup> That Moran holds the offer-of-benefit variant rather than the actual agreement variant is supported by many passages in the text. For example, ‘The epistemic value of his [the testifier's] words is something publicly conferred on them by the speaker’ (2006: 288) suggests that the testifier's assurance by itself gives the addressee an epistemic reason to believe *p*. And again, ‘When an act of telling completes itself, speaker and audience are aligned in this way through their mutual recognition of the speaker's role in determining the kind of reason for belief that is up for acceptance, so that when the speaker is believed there is a non-accidental relation between the reason presented and the reason accepted’ (2006: 299–300). This suggests that an act of telling can be completed regardless of whether the speaker is believed, and it determines the kind of reason for belief, hence gives the addressee an epistemic reason in virtue of the conditional rule of correspondence (that ‘when the speaker is believed . . . accepted’) regardless of whether the addressee accepts the assurance. Despite this, Moran needs to establish the *actual* agreement view to show that reason-giving by assurance involves a personal relation. Hinchman (2005: 572) endorses something similar to an actual agreement view: telling together with the recognition of the intention in telling merely makes an entitlement *available*; the addressee *acquires* an entitlement only by *trusting the teller* (a notion not far from an acceptance generating an agreement).

<sup>18</sup> That assurance entails assertion is supported by the observation that there seems to be a contradiction in saying that I assure you that *p* but I don't assert that *p*. It isn't merely that my saying that I don't assert *p* defeats any assurance I give by saying that I assure you that *p*.

<sup>19</sup> I agree with Moran that in these contexts, one does not normally acquire a testimonially justified belief from assertion. But I do not think that this is because the assertions fall short of assurance. Rather, it is because in persuasion and similar contexts it is assumed that one will not believe the speaker merely because the speaker makes an assertion. Even in such a context one could acquire a testimonially justified belief from an assertion without treating the belief as evidence—or at any rate, without treating it as evidence any more than one must to acquire a testimonially justified belief in a case of assurance.

<sup>20</sup> Moran seems to recognize the latter feature of assertion when he contrasts telling and ‘simply giving expression to what's on your mind’ (280). Of course one can assert a

proposition one doesn't believe. And this prompts the disjunctivist account of assertion: assertion is either an expression of belief or a representation of oneself as believing. For variants of the disjunctivist view, see Williams (2002: 74) and Owens (2006).

<sup>21</sup> For the 'knowing' variant, see e.g. Unger (1975: 250–65), DeRose (1991), and Williamson (2000: 252–5). Moran's catalogue also tells against Brandom's suggestion that asserting commits one 'to being able to vindicate the original claim by showing that one is entitled to make it' (Brandom 1994: 171). (Compare Hinchman's similar criticism of Brandom in Hinchman 2005: 569 n.19.) Brandom's suggestion is not even true of *assurance*.

<sup>22</sup> Owens (2006) gives the example of a secret diary. Such a diary may be full of assertions without my addressing them even to my future self (I might know that I will never read the diary).

<sup>23</sup> Moran (2006: 292, 297–8), and Hinchman (2005: 587). The analogy traces back to J. L. Austin (1946: 100).

<sup>24</sup> Moran seems to think that my promising gives you a non-evidential epistemic reason to believe that I will perform the action. I have no decisive objection to this, but I see no reason to think it's true. Suppose I promise to do A, but you know that I don't often keep my promises. I still owe it to you to do A, but you would be a fool to expect me to do it. Moran could respond to this by saying that my promise gives you a *prima facie* epistemic reason to believe I will do A, although this is defeated by your evidence that I will not do A. But I don't find this response convincing.

<sup>25</sup> I postpone the point that the norm of promising is second-personal or bipolar, while the epistemic norm of testimony is not (Darwall 2006; Owens 2006). Owens argues the point by appeal to overhearer cases: the overhearer gains the same epistemic reason to believe from assurance that the addressee does; so the reason given by testimony is not asymmetrical between second and third person, as required for bipolarity. Here I am discussing only the analogy between promising and assuring.

<sup>26</sup> Thus the term 'accepted promise' is redundant: there is not even a promise unless the promisor's offer is accepted by the promisee.

<sup>27</sup> Compare Lackey (2008: 235).

<sup>28</sup> Might one explain the difference between testifying and promising in this regard by saying that variation in the degree of assurance affects only the evidence assurance supplies, not the non-evidential reason it supplies? Then the quality of the epistemic

reason does not vary with the degree of assurance, once evidential considerations are subtracted. This would restore the analogy with promising. But I can see no ground for this. Plausibly, the quality of the epistemic reason given varies with the degree of assurance even in the primitive case, where this variation cannot be explained by the evidential role of the degree of assurance.

<sup>29</sup> A remark of Moran's may suggest that he takes it that my presenting my utterance as having an epistemic status suffices to give an epistemic reason to believe: 'If, unlike a piece of evidence, the speaker's words have no independent epistemic value as a phenomenon, then how do they acquire the status of a reason to believe something? It seems that this can only be by virtue of the speaker's there and then explicitly presenting his utterance as a reason to believe, with this presentation being accomplished in the act of assertion itself' (288). Moran seems to say here that my presenting my utterance as an epistemic reason to believe  $p$  gives you an epistemic reason to believe  $p$ . But my presenting my utterance as such a reason is by itself epistemically inert—a mere presentation is not an invitation or offer, and it makes no sense to speak of your accepting my presentation. I take it that Moran really means to say that my offer to take responsibility gives you a reason.

<sup>30</sup> I set aside here the objection to the actual agreement view implicit in our discussion of the first disanalogy with promising, that an acceptance, hence an actual agreement, is not necessary to give an epistemic reason. I also set aside the objection implicit in the second disanalogy with promising, that an actual agreement should give a *good* epistemic reason if it gives any reason at all, and this rules out a two-tiered view on which the agreement gives only a reason, not yet a good reason.

<sup>31</sup> The reason you have to trust me can be overridden by reasons other than prudential ones. For example, it can be overridden by a practical reason not to trust me that you acquire from accepting a different invitation to trust. Suppose you accept my invitation to trust me that  $p$ . But then someone invites you to trust her that  $q$ , and you accept that invitation. Now you have commitments and prima facie reasons to trust me that  $p$  and to trust her that  $q$ . But suppose it turns out that it's psychologically impossible for you both to trust me that  $p$  and to trust her that  $q$ . These activities place incompatible demands on your cognitive resources. For example, it might be that merely understanding  $q$  uses resources that crowd out trusting that  $p$ . Note that  $q$  need not be a contrary of  $p$ —it may be any proposition trust in which places incompatible demands on your cognition. Then one or both of the prima facie practical reasons the two invitations give you to trust are overridden. You have an overriding practical reason at most to trust one of us, not to trust each of us. Nothing like this can happen to your epistemic reason to believe. The

psychological impossibility of your trusting me that  $p$  and trusting her that  $q$  does not entail that you do not have an overriding epistemic reason to believe each of these.

<sup>32</sup> Could a proponent of the actual agreement view reply to the objection from differential defeasibility by proposing that the invitation at issue sets special terms for trust: the trust you are invited to show me is conditional only on epistemic considerations? That is, by the terms of the invitation, you are invited to trust me no matter what countervailing practical reasons you may acquire. If you accept this invitation, your reason to trust is vulnerable only to defeat by epistemic reasons, not by practical reasons—as desired to avoid the objection. In reply, I observe that we never do make such invitations. No sensible person would accept an invitation to trust on the condition of ignoring countervailing practical considerations.

<sup>33</sup> But perhaps not enough for a guarantee that  $p$  or for an epistemic reason to believe that  $p$ . Neither my taking epistemic responsibility for asserting  $p$ , nor the epistemic responsibility of my assertion is enough for me to have an epistemic reason to believe  $p$ . For these may obtain simply because I cannot help but believe  $p$ . But plausibly, what is not enough for me to have a reason to believe is also not enough for you to have a reason to believe, and not enough to guarantee that  $p$ .

<sup>34</sup> Should one say that it is the offer or rather the prospect of the benefit that gives you a reason to believe? Suppose you had good reason to doubt that I will follow through on my offer. If this would prevent my assurance from giving you a reason, then it is the prospect and not the offer that gives the reason.

<sup>35</sup> This assurance view differs from the social responsibility view mentioned in note 3. On the latter view, the hearer is justified in virtue of successfully passing epistemic responsibility for the belief that  $p$  to the testifier. This requires only that there be some recognized way of passing responsibility, not that a reason is given or that an offer is made. On the assurance view, a reason is given in virtue of an offer of epistemic responsibility.

<sup>36</sup> One might interpret the second sentence of the quoted passage as saying that my testimony gives you an epistemic reason to believe only in virtue of an actual, completed agreement between us; whereas, on the proposal that promises the best defense of the assurance view, my testimony gives you an epistemic reason to believe in virtue of my offer, because my offer gives you a reason to accept my offer. But the sentence could mean merely that my testimony gives you an epistemic reason to believe in virtue of my

offer and your potential acceptance, and this is consistent with the more promising proposal I would prefer to attribute.

<sup>37</sup> Unlike Lackey (2008: 232–5), I don't think that Moran needs to maintain that when I become responsible for asserting *p*, this transfers all responsibility for your belief from you to me. It may be that I become responsible for your believing *p*; but you are still responsible for believing me as to whether *p*. This seems to be enough to meet Lackey's objection that even though I assume responsibility for my daughter's behaviour while she plays at your house, she may still herself be irresponsible for breaking your vase.

<sup>38</sup> Lackey (2008: 235) objects to Moran's proposal on the ground that I can take responsibility only if I have epistemic competence. I can't say, 'I will take responsibility for the proper functioning of this aircraft', and succeed in doing what I offer to do, unless I am competent to run the aircraft. It's not clear to me whether to take responsibility I must actually be competent or merely that everyone must justifiedly believe me to be so. But the latter is still enough to object to the proposal that taking responsibility is enough to give a reason. However, Moran doesn't make this proposal. Rather, he proposes that in assurance, I make an offer to be responsible for the truth of *p* in a sense that entails only making reparations if *p* turns out to be false, and that this offer gives an epistemic reason. This offer does not commit me to taking responsibility in a way that requires competence. Of course there is a serious question how this could be enough for giving an epistemic reason.

<sup>39</sup> Of course I do not deny that an offer to take future responsibility for a present assertion can be sensible under certain conditions. You may have reason to accept this offer if you don't care whether there is reason to believe *p*, but do want compensation should it turn out that there is no reason to believe it.

<sup>40</sup> See also Owens (2006).

<sup>41</sup> Hinchman (forthcoming: section II) draws the opposite lesson from the overhearer case: that Moran's view does not deliver what is wanted from an assurance view—a difference in the reason given to Kate and Earl—because the relation of assurance and acceptance merely depends on self-presentation and its recognition, and this recognition is available to Earl. Hinchman proposes to remedy this deficiency by making the reason given depend on the testifier's responsiveness to the addressee's epistemic needs. More exactly, the assurance makes the reason available because it is governed by norms that require the responsiveness. Hinchman allows that an unresponsive testifier still gives a



reason in virtue of the fact that the assurance is so governed. I cannot address here this interesting version of the assurance view.

<sup>42</sup> Ross (1986: 79).

<sup>43</sup> The assertion view is quite close to Lackey's (2008: ch. 3) *statement* view of testimony, although her target is testimonial knowledge, not reason-giving. Lackey motivates the statement view with the case of the Creationist Teacher. I think that the teacher must be taken to assert evolutionary theory (despite not believing it) if she is to give a testimonial reason to believe *p*. If we identify statement with assertion, then Lackey's reliabilist statement view entails the assertion view (on the plausible assumption that testimonial knowledge-transfer entails testimonial reason-giving). The assertion view is also close to Owens' (2006) *belief expression* model of testimony. On this view, sincere assertion transmits a rational belief to the hearer: if I rationally believe *p*, my testimony that *p* gives you a reason to believe *p* by expressing my belief that *p*, provided that you trust what I say by accepting it. (Owens ties this to the social responsibility view by taking a transmission of rational belief to depend on passing epistemic responsibility from the hearer to the testifier. See my (2006) for criticisms of the social responsibility view.) Owens denies that an insincere assertion transmits a rational belief. Lackey's Creationist Teacher case shows that this is too strong. Testimony can give a reason to believe *p* without expressing belief. Insincere assertion, if tied to good reasons to believe in the possession of the testifier, is enough to give a good reason to believe. Owens is right to analogize testimony to memory, as I do in (2006), but the analogy does not require belief-transmission, only reason-transmission, for which assertion is enough.

<sup>44</sup> I suppose we should concede to Moran that in these non-assurance cases, the hearer is given no reason to believe the testifier, only to hold the belief. But this can mean only that in these cases the testimony gives no practical reason to believe (aiming at believing what is truly or responsibly asserted). And this difference, as I have argued, is not epistemically significant.

<sup>45</sup> Should we seek to rule out cases of double-bluffing as giving non-reductive epistemic reasons? Suppose I intend to produce in you the belief that I am going to Minsk by anticipating that you mistrust me enough to reason that I am trying to deceive you into thinking that I am not going to Minsk, and this causes you to believe that I am going to Minsk. This is a case of assertion and, in some sense, the hearer's believing because the testifier asserts. It is not a case of assurance (Moran 2006: 290). But I am inclined to think that in the double-bluffing case I do give you a non-evidential testimonial epistemic reason. Of course what justifies your belief is not this non-evidential reason but the

evidential reason from which you elaborately reason to your belief. That is to say, the non-evidential reason merely *propositionally* justifies you in believing the proposition, while the evidential reason *doxastically* justifies your belief. I thank Adam Leite for bringing the relevance of the double-bluffing case to my attention.

<sup>46</sup> The assertion view does not hold, as Moran does, that in any non-evidential testimonial reason-giving, the hearer believes the testifier in the sense singled out by Moran. Denying this is a consequence of covering the primitive cases. So any view that covers those cases will deny it.

<sup>47</sup> I would like to thank Ted Hinchman, Adam Leite, and two anonymous referees for helpful comments on this chapter.

## 12 The Epistemology of Silence

<sup>1</sup> Below I will note that the relevant reliability is not the *reporting-reliability* of the source (as when most of the sources' reports would be true), so much as it is the *coverage-reliability* of the source (as when most of the relevant truths would be reported).

<sup>2</sup> What should be said of memory's preserving what was originally a testimonial belief? I think these should be treated as compound cases, where one's epistemic reliance is on another subject as information source, but where the preservation of the information is a joint effort including both the testifier and the subject. The preservative role of the testifier is this: she must preserve a content she originally acquired, and perform a speech act that enables her audience to quickly and easily recover the content of the information itself. That the testifier has a role in content preservation here should not obscure that the subject herself also has such a role: the subject is relying on her own reliable apprehension of the testimony that confronted her. As argued in Goldberg (2007b), this reliability can be analysed as involving (i) reliability in discerning the force and content of the speech act (reliable comprehension), and (ii) reliability in discriminating reliable from unreliable testimony (reliable filtering).

<sup>3</sup> It is to Burge (1993) that I owe my use of 'content preservation', as covering both memory (or at least what Burge calls 'preservative' memory) and testimony.

<sup>4</sup> Indeed, it is precisely this point that motivates Goldman's well-known (1986) distinction between belief-dependent and belief-independent processes.

<sup>5</sup> This point, asserting the possibility of failing to come across testimony proffered by a source we (standardly) rely on, is insignificant from the perspective of the usual issues that come up in the epistemology of testimony. For the reliability of a source's testimony is independent of how likely it is that her intended audience will come across that testimony. However, if her intended audience is relying on her for *coverage*, the likelihood that this audience will fail to come across her testimony is all-important. See Goldberg (forthcoming a) for further discussion.

<sup>6</sup> Thanks to Jesper Kallestrup for this suggestion.

<sup>7</sup> I introduced these in Goldberg (forthcoming a).

<sup>8</sup> I am assuming something like a reliabilist picture here. See Goldberg (2007b) for an extended defence of this picture.

<sup>9</sup> Where  $\phi$  is some source on which subject  $S$  relies throughout interval  $i$  for coverage in domain  $D$ , we can say that  $\phi$  is 'silent' regarding  $p$  (some proposition in  $D$ ) throughout  $i$  when throughout  $i$   $\phi$  has not issued any report that bears directly on the question whether  $p$ . (To be sure, in the standard cases it is not the actual silence of a source, so much as it is the subject's belief in that silence, that will be relevant to an assessment of subject's beliefs formed on the basis of the (alleged) silence of a given source.) We have here the core of the issue before us: when you believe that Jones did not regularly wear any bright red suit, or that to date WMDs have not been found in Iraq, what is the epistemic bearing on these beliefs of the (alleged) 'silence' of a source on whom (or which) you relied for reliable coverage?

<sup>10</sup> The restriction to 'relevant' truths makes this realistic. After all, even as venerable a daily paper as the *New York Times* only promises to publish 'All of the News *That's Fit to Print*'.

<sup>11</sup> In saying this I do not mean that testimony should be treated as evidence; that would be a mistake for reasons developed in Goldberg (2006).

<sup>12</sup> My nomenclature for claims is as follows: claims will be designated as '(X-Y)', where the term substituting for 'X' will tell which sort of reliance is in question ('S' for reliance on a Source, 'P' for reliance on the part of the system that Preserves contents acquired by the source), and the letter substituting for 'Y' will tell what type of case it is ('M' for Memory case, 'T' for Testimony case).

<sup>13</sup> To a first approximation, these cases might be analysed as follows. A belief is acquired at  $t$ ; this acquisition causes a modification in one's memory system (the belief is put in memory; perhaps there is a process involving transmission first to short-term memory, then to long-term memory); and this modification in the memory system is causally responsible for the subsequent sustainment of a belief with a content distinct from the original belief. Presumably, the sustained content is related to the original content in such a way as not to elicit any suspicion of a content-shift on the part of the subject herself. (It is noteworthy that some content externalists have theories of the semantics of memory on which content-shifts of this sort are to be expected under certain conditions. See Ludlow (1995), Gibbons (1996), and Goldberg (2007c) for a discussion.)

<sup>14</sup> While such a possibility is clear in the abstract, it would appear to be a substantial empirical question whether, and if so to what extent, human memory is susceptible to such failures. Suppose  $\langle p \rangle$  is a proposition that is originally the content of a belief we acquired at some earlier time, but where memory shifts the content of this belief, so that on recollection we have the memory impression that  $q$  (where  $\langle q \rangle$  and  $\langle p \rangle$  are logically independent of one another). Then, when at some later time one considers whether  $p$ , it may be that, despite the fact that one presently has no recollection of  $\langle p \rangle$ , one still has the 'sense' that it would be wrong to conclude that  $\sim p$  on this basis. So even if we do suffer here from a failure of content-preservation, our so suffering would not incline us to form the corresponding 'silence'-based belief.

<sup>15</sup> Here I note that (S-T) does not identify any particular source; it speaks only of 'one of the sources on whom I rely'.

<sup>16</sup> In this respect the motto of the *New York Times* is noteworthy for its coverage-related ambition: 'All of the News That's Fit to Print'.

<sup>17</sup> See Goldberg (forthcoming b) for a full presentation of this idea.

<sup>18</sup> Of course, a justified silence-supported belief would also require S to be sensitive to the likelihood that S herself would come across any relevant report were one to be made. See Goldberg (forthcoming a).

<sup>19</sup> See Goldberg (2007a: ch. 5) for details.

<sup>20</sup> See e.g. Burge (1993), Dummett (1994), Christensen and Kornblith (1997), Lipton (1998), Edwards (2003), Lackey (2005), and Goldberg and Henderson (2006).

<sup>21</sup> I develop this claim at length in Goldberg (forthcoming b).

<sup>22</sup> To make the case do the job I want it to do, we have to imagine that any relevant facts would only be reported by the news sources on which you rely. No doubt, this assumption is unrealistic. But cases are easily imaginable where such an assumption would be realistic: suppose the topic is one where there are only few people interested in the topic, and only one news source that covers such a topic, etc. I will disregard the unrealistic nature of my example in what follows.

<sup>23</sup> This is a point made in Lehrer (1997) as well as in Foley (2006).

<sup>24</sup> Here, and in the suggestions that follow, I am indebted to Alan Millar for very helpful comments.

<sup>25</sup> I should perhaps acknowledge that, while I find the hypothesis of extended cognition to be an intriguing one, I am dubious about its truth. I offer the present argument to those who do not share my scepticism on this score.

<sup>26</sup> That is, other people are not such repositories unless the subject herself testified to them, and they go on to store the subject's information. This is not the sort of case we are envisaging here.

<sup>27</sup> Indeed, so far as I can tell, the epistemological issues surrounding the notion of *newsworthiness* have received no attention whatsoever from epistemologists (as opposed to folks in journalism schools). I regard this as a lacuna in the epistemology literature. I hope to address this at some future time.

<sup>28</sup> I would like to thank Frank Döring for suggesting to me the interest (for social epistemology) of cases like the ones considered here—where you reason that if such-and-such were true, you would have heard about it by now. I would also like to thank the members in the audience at a conference on Social Epistemology at Stirling University, where I presented a paper on a similar theme: Jessica Brown, Igor Douven, Miranda Fricker, Alvin Goldman, Peter Graham, Klemens Kappell, Jon Kvanvig, Jennifer Lackey, Peter Lipton, Nenad Miscevic, Alan Millar, Erik Olsson, Duncan Pritchard, Ernie Sosa, and Finn Spicer, for particular comments on that paper. Special thanks to Jasper Kallestrup, who served as my commentator at that conference, and whose suggestions inform the present chapter in a variety of ways; and to Alan Millar, for extensive comments on an earlier draft of this chapter.

### 13 Epistemic Circularity and Epistemic Incommensurability

<sup>1</sup> See Alston (1989), Sosa (1997a, 1997b). For a novel defence of this strategy that links externalism to Reidian common-sensism, see Bergmann (2004). Reed (2006) offers a criticism of Bergmann's position.

<sup>2</sup> See e.g. Fumerton (1995), Stroud (1994), and Pritchard (2005).

<sup>3</sup> This is one of the lessons that can be drawn from S. Cohen's (2002) well-known discussion in 'Basic Knowledge and the Problem of Easy Knowledge'. Cohen himself is targeting the notion of basic knowledge in general.

<sup>4</sup> I have defended the value of truth elsewhere, see Lynch (2004, 2009a, 2009b).

<sup>5</sup> This illustrates a general lesson: where the methods in question each claim privileged access to a special domain of facts, appeals to other principles will be apt to lead to non-inferential principles that are themselves subject to epistemic circularity. Consider, to take another example, where Hylas and Philo disagree about the only reliable method for determining what is morally good. Hylas claims it is divine revelation; Philo rational intuition. It is hard to see how mutual circularity can be avoided in such a case, since both principles claim that a certain non-inferential method has privileged access to a restricted domain of facts. Any other method will be rejected as not being capable of accessing those facts.

<sup>6</sup> For this reason, deep epistemic disagreements are not properly classed as what some have come to call 'reasonable disagreements'. A reasonable disagreement in this sense is one where each side recognizes that they are drawing equally reasonable but incompatible conclusions from shared evidence. Deep epistemic disagreements don't fit this model. For I can't say of someone who does not believe my FEP that we 'draw distinct conclusions from shared evidence' for the simple reason that I won't believe that we share evidence if we don't share *what counts as evidence*. Think of a situation where an advocate of divine revelation tells me that *p*. I doubt that *p* is the case because I doubt that reliability of divine revelation. I may claim we have the same evidence, my opponent will not.

<sup>7</sup> Note the 'probably' in the conclusion. The argument is deductive, but its conclusion inherits its content from the abduction in the second premise.

<sup>8</sup> Compare Thomas Nagel: 'To reason is to think systematically in ways anyone looking over my shoulder ought to be able to recognize as correct' (1997: 5). Here Nagel oversteps: he suggests that reasoning itself (and therefore possessing a reason presumably) requires that the reason so possessed be recognizable by anyone as a reason.

Not so: I can have a belief about my own mental states that may be a reason for some other belief that could not possibly be recognized by you as a reason for the latter belief. Nonetheless, the point holds where we are concerned with giving a reason, as opposed to being or having a reason.

<sup>9</sup> It would be like telling someone faced with a moral dilemma, where e.g. the demands of utility clash with respect for persons, not to worry because, ‘if you act on the correct moral principles, you’ll do the right thing’. Thanks a lot.

<sup>10</sup> For a non-factualist account, see H. Field (2001). Steven Hales (2006) discusses a relativist solution to a similar problem.

<sup>11</sup> See Pritchard (2005: 245–9).

<sup>12</sup> What I am calling the epistemic method game was first suggested to me by Eberhard Herrmann; I am not sure he would approve of its use here, however.

<sup>13</sup> This chapter has benefited from discussions at the University of St Andrews, the University of Edinburgh, the Bled Epistemology Conference, and the University of Connecticut. My thanks to Don Baxter, Tom Bontly, Matthew Chrisman, Patrick Greenough, Terry Horgan, Scott Lehman, Alan Millar, Baron Reed, and, especially, Duncan Pritchard for additional comments.

## 14 The Epistemology of Disagreement

<sup>1</sup> Or, perhaps more generally, what they need is not *reason* but *justification*, or, more generally still, *positive epistemic status*; justification and reasons may perhaps provide only some among the ways to gain such status. This footnote is intended to forestall misinterpretation, allowing the main text latitude to proceed with the simplifying assumption that ‘reason’ is required, although we shall eventually find reason to reject this as an oversimplification.

<sup>2</sup> Feldman (2007) adopts the view under that label; see also Christensen (2007) and White (2005) for additional support. Strictly speaking, it is not a body of evidence *E* that justifies withholding. What matters is, rather, the *fact* that *E* is *S*’s *total* evidence bearing on  $\langle p \rangle$ ; it is this fact that might justify *S*’s withholding on  $\langle p \rangle$ . Thus arises the question of how to characterize the abstract relation of justifying so that it can apply equally in both

cases, both when it is believing that is justified, and when it is withholding. But we won't tarry over this.

<sup>3</sup> Suppose Pro does switch from believing to withholding, while Con holds his ground, whereupon they disclose once again. What now is required of Pro? Plausibly, if neither of them has acquired any independent reason to downgrade the other, each should take into account the stance of the other, and the two should jointly settle on some intermediate position. So, what is a position intermediate between Pro's withholding and Con's disbelief? If Con now moves from disbelieving to withholding, this is just to join the opposition, which seems wrong if they should really be splitting the difference. *Is there some position intermediate between withholding and disbelieving?* Surprisingly enough, it can be argued that there is: namely, deliberately abstaining *both* from positively forbearing and *also* from disbelieving. Forbearing involves more, after all, than just 'not believing'. It is something more positive than that: namely, deliberately bringing it about that one fails to believe. Is this now what they are both required to do: to *abstain* both from disbelieving and from withholding? To occupy this position, then, one must deliberately refrain from deliberately bringing about each (severally) of two outcomes: that one believes  $\langle \text{not-}p \rangle$ , and that one fails to believe  $\langle \text{not-}p \rangle$ . Of course, one must either believe or not believe. But one need not *bring about* either one of these. One might just deliberately abstain from both options. Unfortunately, a closer look makes it doubtful that there is any such coherent intermediate position. After all, if one deliberately abstains from disbelieving and deliberately abstains from believing, that is tantamount to positively forbearing. So one cannot abstain from both disbelieving and forbearing, except by positively believing.

If there is no such coherent position intermediate between withholding and disbelieving, however, consider then the options open to our protagonists: to withholder Non and disbeliever Con. What are they to do upon taking note of their disagreement given that neither can downgrade the other, provided that in their situation believing is not an option?

Here I will not press further along that path. Nor will I switch to a conception that incorporates degrees of belief between 0 and 1, or even replaces the threefold structure of attitudes—belief, withholding, and disbelief—with a spectrum of degrees of confidence between 0 and 1. Let us shelve such issues in favor of a focus on the stronger disagreement where one of the parties affirms what the other denies. Can they both be reasonable?



<sup>4</sup> It might be argued that when a controversy is hoary, the key terms involved cannot have different senses or meanings or proper interpretations, so that the participants could be speaking past each other. And this is allegedly because the terms acquire their proper senses, etc., from the aim of the disputants to participate in the traditional dispute. This I do not find nearly as plausible as William James's view of a dispute by a campfire as to whether a man goes 'around' the squirrel who scurries around the tree always keeping the tree between them as the man runs around full circle. The man goes first to the north, then to the east, then to the south, and then to the west, and then back to the north of the squirrel, so he goes around the squirrel in one sense. However, he fails to go around the squirrel by failing to be first in front, then on one side, then in back, then on the other side, and then back in front. Plausibly there are these two senses of 'going around', two senses which no amount of firm intention to clash among disputants would seem to remove. This approach fits also Hume's remark in his *Enquiry*, Section 8 that '[f]rom this circumstance alone, that a controversy has been long kept on foot, and remains still undecided, we may presume, that there is some ambiguity in the expression, and that the disputants affix different ideas to the terms employed in the controversy' (Hume 2000: 62).

<sup>5</sup> See Christensen (2007), Elga (2007), and Feldman (2006). For relevant disagreement, see Kelly (2005), Rosen (2001), and van Inwagen (1996). I will side with the opposing camp for the most part, though my reasons in what follows are multiply different from those to be found in the earlier literature. Elga explicitly relates his rejection of downgrading without independent reason to a certain well-known objection to reliabilism, and he cites some of my own work in this connection. I agree that the comparison is apt, and myself take the same position on both sides of it. In the earlier work I defend the position as a solution to the Pyrrhonian problematic. In what follows I defend the same position (or a close neighbour) on the epistemology of disagreement.

<sup>6</sup> The restaurant and watch examples are Christensen's, the quad example Feldman's.

<sup>7</sup> Here are some examples:

- a. Epistemic peers are equals in regard to 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (Gutting 1982: 83).
- b. Epistemic peers with respect to a question are 'equals in regard to their familiarity with the evidence and arguments which bear on that question, and . . . equals with respect to general epistemic virtues such as intelligence, thoroughness, and freedom from bias' (Kelly 2005).

- a. Epistemic peers are equals in regard to 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (Gutting 1982: 83).
- c. X and Y are epistemic peers with respect to a certain question if and only if conditional on a disagreement arising between them on that question the two of them are equally likely to be mistaken (Elga 2007). Elga insists that peerage on a specific question at a specific time between two subjects, X and Y, requires his condition. For if one of them is more likely to be right on that question at that time, then the two are not peers *on that question*, regardless of how generally virtuous or specifically well informed that subject may be. But how then are we to understand such 'likelihood'? For example, if it is likelihood relative to the information and virtues possessed by the thinker, then the two definitions, Elga's and Kelly's, collapse into one. (Actually, we may need to add to Kelly's definition that the two subjects not only *possess* those virtues, but *employ* them equally as well. So let's take this to be implicit in that definition.)

<sup>8</sup> We consider this question as posed through the present definition of peerage, even if we might *also* find interest in parallel questions that use definitions less restricted to the specific locus of the disagreement.

<sup>9</sup> Elga's position is pleasingly clear and unambiguous: 'When you learn of your friend's opposing judgment, you should think that the two of you are equally likely to be correct . . . . If it were reasonable for you to give your own evaluation extra weight—if it were reasonable to be more than 50% confident that you are right—then you would have gotten some evidence that you are a better evaluator than your friend. But that is absurd' (Elga 2007).

<sup>10</sup> Feldman argues persuasively that this sad outcome is to a large extent our fate.

<sup>11</sup> Here in brief is part of the view, as stated by Elga: '[With] respect to many controversial issues, the associates who one counts as peers tend to have views that are similar to one's own. That is why — contrary to initial impressions — the equal weight view does not require one to suspend judgment on everything controversial' (Elga 2007).

<sup>12</sup> 'Rather than taking her views on the surrounding issues for granted, Ann should attend to the larger disagreement between her and Beth: disagreement about a whole cluster of issues linked to abortion' (Elga 2007). Also: '[Setting] . . . aside her reasoning about the issues in the cluster, and setting aside Beth's opinions about those issues, Ann does not think Beth would be just as likely as her to get things right. That is because there is no fact of the matter about Ann's opinion of Beth, once so many of Ann's considerations

have been set aside. Hence the equal weight view does not require Ann to suspend judgment about the cluster. That blocks the objection' (Elga 2007).

<sup>13</sup> Perhaps it will be contended that the relevant likelihood estimates are relativized. So it might be held that what follows from our commitment to our likelihood of being right, from *our* point of view, is the unlikelihood *from our point of view*, of someone else being right if they contradict us. But that, it may be added, is trivial. 'The relevant issue for us,' it may then be concluded, 'the relevant more interesting issue, is that of whether our opponent is likely to be right from *their* point of view'. Perhaps. But if the relativization here is to the evidence and competence of the subject, then if deep and broad disagreement is compatible with different likelihood profiles for propositions under dispute, so that one and the same proposition might be likely relative to S and unlikely relative to S', then it must be that S and S' do not *fully* share their relevant evidence and competence. So this would not be the sort of case that concerns us here at this stage.

<sup>14</sup> According to Feldman (2006: 232) '[p]erhaps the evidential relation can hold in the absence of a believer's realization that it obtains. But once one thinks explicitly about the topic, it becomes harder to see how it can remain reasonable to maintain belief once one realizes that one has no good reason to think one's evidence supports that belief. In effect, this realization serves as a defeater of whatever support the original evidence provided.'

<sup>15</sup> Here I have assumed that there are at least *some* possible cases of such disagreement, where someone denies what to us is quite obvious. It might be thought that simple arithmetic could never provide an example, but consider an adult learning arithmetic who has nearly learned his arithmetical tables, but not quite. Can he not believe that, say,  $6+7=14$ ? Prior to delivering himself of that opinion, we might have held him a peer, if we basically knew him as just a fellow adult. Indeed, independently of his unfortunate belief, we might have no reason whatever to hold ourselves superior to him on matters of arithmetic. Indeed, he might have been our dining companion in the original restaurant example. So, if we now downgrade him it can only be based on the substance of the disagreement, or so it would seem, in which case we have here a case of proper downgrading based on the substance of a disagreement, even without any independent reason to consider our opponent inferior.

<sup>16</sup> 'I agree . . . that if I don't know now that I'm not dreaming, it follows that I don't *know* that I am standing up, even if I both actually am and think that I am. But this first part of the argument is a consideration which cuts both ways. For, if it is true, it follows that it is also true that if I *do* know that I am standing up, then I do know that I am not dreaming. I can therefore just as well argue: since I do know that I'm standing up, it follows that I do

know that I'm not dreaming; as my opponent can argue: since you don't know that you're not dreaming, it follows that you don't know that you're standing up. The one argument is just as good as the other, unless my opponent can give better reasons for asserting that I don't know that I'm not dreaming, than I can give for asserting that I do know that I am standing up' (Moore 1959a: 247).

<sup>17</sup> 'How am I to prove now that 'Here's one hand and here's another'? I do not believe I can do it. In order to do it, I should need to prove for one thing, as Descartes pointed out, that I am not now dreaming. But how can I prove that I am not? I have, no doubt, conclusive reasons for asserting that I am not now dreaming; I have conclusive evidence that I am awake; but that is a very different thing from being able to prove it. I could not tell you what all my evidence is; and I should require to do this at least, in order to give you a proof' (Moore 1959b: 149).

<sup>18</sup> From J. L. Austin underline (1962: 48–9):

I may have the experience (dubbed 'delusive' presumably) of dreaming that I am being presented to the Pope. Could it be seriously suggested that having this dream is 'qualitatively indistinguishable' from *actually being* presented to the Pope? Quite obviously not. After all, we have the phrase 'a dream-like quality'; some waking experiences are said to have this dream-like quality, and some artists and writers occasionally try to impart it, usually with scant success, to their works. But of course, if the fact here alleged *were* a fact, the phrase would be perfectly meaningless, because applicable to everything. If dreams were not 'qualitatively' different from waking experiences, then *every* waking experience would be like a dream; the dream-like quality would be, not difficult to capture, but impossible to avoid. It is true . . . that dreams are *narrated* in the same terms as waking experiences: these terms, after all, are the best terms we have; but it would be wildly wrong to conclude from this that what is narrated in the two cases is *exactly alike*. When we are hit on the head we sometimes say that we 'see stars'; but for all that, seeing stars when you are hit on the head is *not* 'qualitatively' indistinguishable from seeing stars when you look at the sky.

Compare also the last paragraph of Descartes's *Meditations*:

I know that in matters regarding the well-being of the body, all my senses report the truth much more frequently than not. Also, I can almost always make use of more than one sense to investigate the same thing; and in addition, I can use both my memory, which connects present experiences with preceding ones, and my intellect, which has by now examined all the causes of error. Accordingly, I should not have any further fears about

the falsity of what my senses tell me every day; on the contrary, the exaggerated doubts of the last few days should be dismissed as laughable. This applies especially to the principal reason for doubt, namely my inability to distinguish between being asleep and being awake. For I now notice that there is a vast difference between the two, in that dreams are never linked by memory with all the other actions of life as waking experiences are. If, while I am awake, anyone were suddenly to appear to me and then disappear immediately, as happens in sleep, so that I could not see where he had come from or where he had gone to, it would not be unreasonable for me to judge that he was a ghost, or a vision created in my brain [. . .] like those that are formed in the brain when I sleep; (added in the French version), rather than a real man. But when I distinctly see where things come from and where and when they come to me, and when I can connect my perceptions of them with the whole of the rest of my life without a break, then I am quite certain that when I encounter these things I am not asleep but awake. And I ought not to have even the slightest doubt of their reality if, after calling upon all the senses as well as my memory and my intellect in order to check them, I receive no conflicting reports from any of these sources. For from the fact that God is not a deceiver it follows that in cases like these I am completely free from error. But since the pressure of things to be done does not always allow us to stop and make such a meticulous check, it must be admitted that in this human life we are often liable to make mistakes about particular things, and we must acknowledge the weakness of our nature. (Descartes 1984: 61–2)

<sup>19</sup> An appendix takes up more fully the nature of Moore's disagreement with the sceptic.

<sup>20</sup> Indeed, he takes his case to be worse than that: he thinks we can be sure that we *cannot* fully expound our evidence. But our weaker claim holds nonetheless: his case is one where we are unsure that we have fully expounded our evidence.

<sup>21</sup> Presented at the Stirling Social Epistemology conference of August 2007, at a Brown workshop of February 2008, and a Princeton workshop of April 2008, with helpful discussion following on each occasion.

<sup>22</sup> The proposition on which they disagree is essential to Moore's *argument* that he knows he is not dreaming, though he would not likely regard any such argument as essential to his *knowledge* that he is not dreaming, or perhaps even to his knowledge that he knows he is not dreaming. A fuller treatment of Moore's epistemology constitutes the first chapter of my *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Vol. II (Sosa (2009)).

## 15 A Justificationist View of Disagreement's Epistemic Significance

<sup>1</sup> Of course, I am not really biased in my holding this belief!

<sup>2</sup> Much will depend on which epistemic asymmetries are here relevant; I shall address this issue in detail later in this chapter.

<sup>3</sup> I borrow this term from Thomas Kelly (2005) who, in turn, borrows it from Gutting (1982). I shall later give a more precise characterization of what is involved in being epistemic peers.

<sup>4</sup> See also van Inwagen (1996).

<sup>5</sup> When I speak merely of 'peer disagreement' I mean disagreement between *epistemic* peers.

<sup>6</sup> I borrow the phrase 'extra weight' from Elga (2007).

<sup>7</sup> See also Wedgwood (2007: 262).

<sup>8</sup> There are passages in Fumerton (2010) that also echo egocentric nonconformism.

<sup>9</sup> For a different version of nonconformism, one that combines a principle of epistemic conservatism with an appeal to the underdetermination of theory by evidence, see Moffett (2007).

<sup>10</sup> In other words, the conformist claims that one cannot downgrade the epistemic status of one's peer merely because of the disagreement in question itself; one must have an independent reason for so downgrading. Thus, Adam Elga writes, 'Suppose that . . . you and your friend are to judge the truth of a claim, based on the same batch of evidence . . . . *Without some antecedent reason to think that you are a better evaluator, the disagreements between you and your friend are no evidence that she has made most of the mistakes*' (Elga 2007: 487, emphasis added).

<sup>11</sup> Full disclosure, according to Feldman, occurs when the parties to the disagreement 'have thoroughly discussed the issues. They know each other's reasons and arguments, and that the other person has come to a competing conclusion after examining the same information' (Feldman 2006: 220).

<sup>12</sup> In particular, he argues that when faced with peer disagreement, '(1) I should assess explanations for the disagreement in a way that's independent of my reasoning on the

matter under dispute, and (2) to the extent that this sort of assessment provides reason for me to think that the explanation in terms of my own error is as good as that in terms of my friend's error, I should move my belief toward my friend's' (Christensen 2007: 199).

<sup>13</sup> Even more precisely, Elga writes: 'Equal weight view: Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement' (Elga 2007: 490).

<sup>14</sup> A possible exception is van Inwagen (1996, 2010), who focuses specifically on disagreements that are 'matters of interminable debate' (1996: 141), such as those found in philosophy, politics, and religion. I should mention, however, that van Inwagen fails to offer any general arguments or principles to justify treating various kinds of disagreement differently.

<sup>15</sup> I borrow this term from Christensen (2007).

<sup>16</sup> Christensen calls this 'Cognitive Parity'.

<sup>17</sup> Adam Elga offers an alternative characterization, according to which you count someone as an epistemic peer if 'you think that, conditional on a disagreement arising, the two of you are equally likely to be mistaken' (2007: 487). Elga defends his 'non-standard' usage of 'epistemic peer' on the following grounds:

On more standard usages, an epistemic peer is defined to be an equal with respect to such factors as 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (Gutting 1982: 83), 'familiarity with the evidence and arguments which bear on [the relevant] question', and 'general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias' (Kelly 2005). In defense of my use, suppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you are to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, at least on that occasion. (Elga 2007: 499 n. 21)

Now, even if Elga's criticism of the standard use of 'epistemic peer' is correct, it would not apply to my characterization in the text since I specify that A and B are both evidential and cognitive equals *relative to the question whether p*. Thus, according to my

use of ‘epistemic peer’, two people could not be evidential and cognitive equals with respect to the question whether  $p$  and yet deviate in their likelihood to be mistaken regarding this question.

There are, however, independent reasons to question Elga's non-standard use of this term. For on his account, two people could radically differ in both their evidential backgrounds and their cognitive abilities with respect to the question whether  $p$ , yet nonetheless turn out to be epistemic peers regarding this question. For instance, I may be a complete novice with respect to identifying birds of prey, and you may be an expert ornithologist. When I am sober and you are highly intoxicated, however, we may be equally likely to be mistaken about whether the bird flying overhead is an osprey. On Elga's account, then, you and I would be epistemic peers with respect to this question, but this strikes me as quite a counter-intuitive result.

<sup>18</sup> It is unclear whether full disclosure for Feldman includes that A and B both *know* that all of the evidence relevant to the question whether  $p$  has been shared, but I will assume that this is the case at this point. Later in this chapter, I will consider the epistemic significance of less idealized forms of disagreement.

<sup>19</sup> More precisely, this condition (and the one to follow in my characterization of ordinary disagreement) should be expressed as follows:

(2\*) prior to recognizing that this is so, A and B would take themselves to be epistemic peers with respect to this question, were they presented with the relevant concepts and definitions involved in being such peers.

This version of the condition makes clear that the relevant parties need not literally possess the concepts of evidential equality, cognitive equality, and full disclosure in order to participate in a disagreement; it is sufficient that such parties would regard one another as epistemic peers, were they in possession of these concepts. For ease of expression, however, I shall stick with the less cumbersome formulation found in (2), leaving the reader to interpret it along the lines found in (2\*).

<sup>20</sup> Otherwise put, A and B disagree in an ordinary sense if and only if, relative to the question whether  $p$ , (1) A and B are aware that they hold differing doxastic attitudes, and (2) prior to recognizing that this is so, A and B are not aware of any relevant epistemic asymmetry between their situations.

<sup>21</sup> What about cases where the parties to the debate *should* take one another to be epistemic peers, but do not? Such cases may raise all sorts of interesting epistemological



questions about egoism, dogmatism, irrationality, and the like. But these issues seem rather independent of the epistemic significance of disagreement itself, and thus lie outside the scope of this chapter.

<sup>22</sup> This may lead one to wonder whether there is a certain amount of talking past one another taking place in this debate.

<sup>23</sup> At the end of this section, I shall argue that this assumption is, in fact, incorrect.

<sup>24</sup> For instance, recall that Rosen says, 'It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence . . . it would appear to be a fact of epistemic life that a careful review of the evidence does not guarantee consensus, even among thoughtful and otherwise rational investigators' (Rosen 2001: 71–2).

<sup>25</sup> Recall, for instance, that Feldman writes, 'in situations of full disclosure, where there are not evident asymmetries, the parties to the disagreement would be reasonable in suspending judgement about the matter at hand. There are, in other words, no reasonable disagreements after full disclosure' (Feldman 2006: 235).

<sup>26</sup> It is of interest to note that proponents of conformism often appeal to analogies with non-agential instruments in an attempt to show the absurdity of nonconformist results. For instance, Christensen provides the following example:

I look at my watch, a one-year-old Acme that has worked fine so far, and see that it says 4:10. Simultaneously, however, my friend consults her watch—also a one-year-old Acme with a fine track record—and it reads 4:20. When she tells me this, it clearly gives me new evidence that her watch is fast: I should not trust her watch as much as I would have before finding out that it disagreed with mine. But just as clearly, I've just gotten new evidence that my watch is slow, and this should diminish my trust in it. In this case, it's obvious that the fact that one of the watches is on my wrist does not introduce an epistemically relevant asymmetry. (2007: 196)

Feldman (2006: 234) makes a similar point appealing to different thermometers. The intuition we are invited to share by conformists here is that, without a reason independent of the disagreement in question for doing so, it is just as absurd to prefer my belief over my peer's as it is to prefer the time my watch says over that of my friend's. But notice: there are cases analogous to PERCEPTION involving non-agential instruments. Consider the following:

WATCH: Sonya and I eat lunch together at a restaurant at noon, take a long walk around the lake, shop at multiple stores, and read at the bookstore café for several hours. I look out the window, noticing that the sun is setting, and say, ‘It is 7:45 PM, so we should get going’, after which Sonya responds, ‘My watch says it is only 1:15 PM’.

It seems clear to me in WATCH that, despite not having a reason independent of the disagreement to prefer the time my watch says over that of Sonya's, I am rational in regarding hers as the inaccurate one. For it is simply not at all plausible to think that we could have eaten lunch, walked around the lake, gone shopping, and read for several hours at the bookstore café in an hour and fifteen minutes, nor is it likely that the sun would be setting in Chicago at 1:15 in the afternoon. Thus, I take it that conformists fail to garner the intuitive mileage that they hope to on behalf of their view from these sorts of non-agential cases.

<sup>27</sup> I should mention that Christensen (2007) and Elga (2007) consider a counterexample to conformism that bears some similarities to those found here and attempt to show how their respective views have the resources to rule out doxastic revision in such a case. In Section 3 of this chapter, I shall consider their responses in some detail and argue that they fail in various respects.

<sup>28</sup> I borrow this phrase from Nathan Christiansen.

<sup>29</sup> Perhaps this is why Kelly claims that ‘to uncritically assume that things are perfectly symmetrical with respect to all of the epistemically relevant considerations . . . is . . . to subtly beg the question in favor of the skeptical view’ (2005: 178–9).

<sup>30</sup> Alvin Goldman makes a similar point in his (2010).

<sup>31</sup> Henceforth, when I speak merely of ‘disagreement’, I mean ordinary disagreement.

<sup>32</sup> See, for instance, Sosa (this volume, ch. 14).

<sup>33</sup> I address this question in more detail in my (forthcoming).

<sup>34</sup> It should be noted that there are two probabilities that need to be kept distinct: there is (i) the subjective probability that  $p$  is true, and (ii) the subjective probability that I am correct in my belief regarding  $p$ . Earlier in the chapter, the discussion focused on (i), but now I am emphasizing (ii). It should be clear, however, that (i) and (ii) are intimately related for the conformist. For instance, recall David Christensen's claim that, upon discovering that I disagree with an epistemic peer, ‘(1) I should assess explanations for the disagreement in a way that's independent of my reasoning on the matter under

dispute, and (2) to the extent that this sort of assessment provides reason for me to think that the explanation in terms of my own error is as good as that in terms of my friend's error, I should move my belief toward my friend's' (Christensen 2007: 199). Here Christensen is saying that in cases of peer disagreement, to the extent that I am willing to assign a 50% probability both to my being correct in my belief regarding  $p$  and to my epistemic peer's being correct, I should split the difference with my epistemic peer in the degree to which I believe that  $p$ . Thus, the probabilities assigned with respect to (ii) directly determine the probabilities that should be assigned with respect to (i).

<sup>35</sup> I am grateful to Benjamin Almassi for this question.

<sup>36</sup> What if the conformist were to respond that cases involving an asymmetry of access to personal information are ones where I in fact have reason, independent of the ordinary disagreement in question, to downgrade my opponent? (I am grateful to David Christensen for this type of response.) Given this, the conformist could grant the intuition that no doxastic revision is required in PERCEPTION, ELEMENTARY MATH, and DIRECTIONS, but then deny that it is a problem for her view. By way of response to this line of thought, notice that such a move has the consequence that there are virtually no actual instances of the kind of disagreement to which the conformist view purportedly applies. For in nearly all cases of ordinary disagreement, there is some sort of asymmetry of personal information. How often, for instance, do I know to the same extent as I do in my own case that my opponent is not distracted, sleep-deprived, melancholy, and so on? Not very often (if ever). But then, given the response under consideration, it would turn out that there are virtually no cases of actual disagreement where equal weight should be given to both my view and my opponent's and, hence, it would turn out that there are virtually no cases where disagreeing peers are rationally required to doxastically conform. This, I take it, is an unwelcome result for conformism.

<sup>37</sup> A further defence of conformism frequently found in the literature appeals to the following:

The Uniqueness Thesis: A body of evidence,  $E$ , justifies at most one doxastic attitude—i.e. believing, disbelieving, suspending judgement—toward any particular proposition.

Indeed, Kelly (forthcoming) argues that 'a commitment to The Equal Weight View carries with it a commitment to The Uniqueness Thesis' (forthcoming: ms, p. 11). For proponents of this thesis, see White (2005), Feldman (2006, 2007), and Christensen (2007). While it lies outside the scope of this chapter to discuss The Uniqueness Thesis, let me say that none of the arguments I make here depends directly on its truth or falsity.

<sup>38</sup> See Christensen (2007: 193). I have slightly modified inessential details of Christensen's case, but all of the elements central to the conclusion about disagreement are the same.

<sup>39</sup> I am assuming that Ramona and I each asserting that we have carefully performed the relevant calculations in our heads suffices for us to have reason to believe that full disclosure has taken place. I should say, however, that the condition of full disclosure itself can be fleshed out in more or less idealized ways. For instance, at one end of the spectrum, it may be required that each of us provides all of the details of our respective calculations to one another. In this case, we again face problems making sense of the possibility of the disagreement. For if we have each gone step by step with one another through our fairly elementary calculations, what room is left for us to continue to disagree about the amount owed? At the other end of the spectrum, full disclosure may require nothing more than each of us asserting that we have arrived at the conclusion in question. But then this does not add anything significant to the initial disagreement. Thus, I think it is best to understand full disclosure as falling somewhere in the middle, which is what I have built into BILL CALCULATION.

<sup>40</sup> By a belief enjoying 'a very high degree of justified confidence' I mean a very confident belief that is highly justified.

<sup>41</sup> In my (forthcoming), I provide a far more detailed defence of my justificationist view that adds crucial further conditions to the above two principles. In particular, I argue that no doxastic revision is necessary if there is either a highly justified target or 'protecting' belief, and that substantial doxastic revision is required if there is neither a highly justified target nor 'protecting' belief. A's belief that  $p$  is protected by A's belief that  $q$  if and only if both A's belief that  $p$  and A's belief that  $q$  are members of a subset of A's beliefs, each of which is challenged by the same instance of ordinary disagreement with B, and where A's belief that  $q$  is highly justified and confidently held.

<sup>42</sup> I am grateful to Nikolaj Jang Pedersen for raising this sort of question to my account.

<sup>43</sup> This reliability condition is certainly compatible with a subjective rationality constraint being necessary for justification as well. For more on this, see my (2008).

<sup>44</sup> Or, at least, roughly equal in terms of reliability or truth-conduciveness.

<sup>45</sup> It should be noted that while Christensen (2007) and Elga (2007) both try to accommodate the intuitions elicited by cases such as PERCEPTION and DIRECTIONS within their conformist framework, Feldman (2006, 2007) instead tries to show that his

view delivers the correct result in such situations: both parties should indeed withhold belief with respect to the question under dispute.

<sup>46</sup> In a similar spirit, Christensen (2007) argues that in order for Mia to have reached such an absurd answer, it is plausible for me to think that she must have failed to use a highly reliable, commonsense method of checking her results. Since I am confident that I did use such a reliable method, I do not have to engage in any doxastic revision in the face of our disagreement.

<sup>47</sup> The only possible exception to this claim is ELEMENTARY MATH, where it may be argued that denying that  $2 + 2$  equals 4 is insane in and of itself. I am inclined to think that this claim of insanity, too, makes sense only against the background of one's extraordinarily justified belief that  $2 + 2$  does equal 4. But even if one does not share my view here, the conformist is still left with the problem of explaining why doxastic revision is not intuitively required in PERCEPTION, DIRECTIONS, and EXTREME BILL CALCULATION in a way that does not appeal to the notion of justification.

<sup>48</sup> For helpful comments on earlier versions of this paper, I am grateful to Larry BonJour, David Christensen, Jeremy Fantl, Richard Feldman, Richard Fumerton, Peter Graham, Jon Kvanvig, Matt McGrath, Nikolaj Jang Pedersen, Blake Roeber, Andrew Rotondo, Ernie Sosa, Peter van Inwagen, audience members at the University of Washington, the Social Epistemology Conference at the University of Stirling, the Disagreement Conference at the University of Calgary, the 2007 meeting of the Eastern APA in Baltimore, Brown University, the University of St Andrews, and, especially, to Baron Reed.

## Bibliography

*Abrams v. U.S.* (1919). 250 U.S. 616.

Adler, J. (1994). 'Testimony, Trust, Knowing', *The Journal of Philosophy*, 91: 264–75.

— (2002). *Belief's Own Ethics*. Cambridge, MA: MIT Press.

Alcoff, L. and Potter, E. (eds.) (1993). *Feminist Epistemologies*. New York: Routledge.

Alston, W. (1985). 'Concepts of Epistemic Justification', *The Monist*, 68: 57–89.

— (1989). *Epistemic Justification*. Ithaca: Cornell University Press.

— (2000). *Illocutionary Acts and Sentence Meaning*. Ithaca: Cornell University Press.

— (2005). *Beyond 'Justification': Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Anderson, E. (1995). 'Feminist Epistemology: An Interpretation and a Defense', *Hypatia: A Journal of Feminist Philosophy*, 10: 50–84.

— (2006). 'The Epistemology of Democracy', *Episteme: A Journal of Social Epistemology*, 3: 8–22.

Annis, D. (1978). 'A Contextualist Theory of Epistemic Justification', *American Philosophy Quarterly*, 15: 213–19.

Anscombe, G. E. M. (1979). 'What is it to Believe Someone?', in C. Delaney (ed.) (1979: 141–51).

Antony, L. (ed.) (2007). *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press.

Ariew, A., Cummins, R., and Perlman, M. (eds.) (2002). *Functions: New Essays in the Philosophy of Psychology and Biology*. Cambridge, MA: MIT Press.

Audi, R. (1988). 'Justification, Truth, and Reliability', *Philosophy and Phenomenological Research*, 49: 1–29.

— (1994). 'Dispositional Beliefs and Dispositions to Believe', *Noûs*, 28: 419–34.

— (1997). 'The Place of Testimony in the Fabric of Knowledge and Justification', *American Philosophical Quarterly*, 34: 404–22.

Austin, J. L. (1946). 'Other Minds', *The Aristotelian Society*, Supplementary Volume, 20: 148–87. Reprinted in J. L. Austin (1979: 44–84). Page references are to the reprinted version.

— (1962). *Sense and Sensibilia*. Oxford: Oxford University Press.

— (1979). *Philosophical Papers*. Oxford: Oxford University Press.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Ayer, A. J. (1946). *Language, Truth and Logic*. London: V. Gollancz.

Bach, K. (2006). 'Speech Acts and Pragmatics', in Devitt and Hanley (eds.) (2006: 147–67).

—— and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.

Barnes, B. and Bloor, D. (1982). 'Relativism, Rationalism, and the Sociology of Knowledge', in Hollis and Lukes (eds.) (1982: 21–47).

Bar-On, D. and Chrisman, M. (2009). 'Ethical Neo-Expressivism', in Shafer-Landau (2009: 133–65).

326

Barth, E. M. and Krabbe E. C. W. (1982). *From Axiom to Dialogue: A Philosophical Study of Logic and Argumentation*. Berlin: de Gruyter.

Beanblossom, R. and Lehrer, K. (eds.) (1983). *Thomas Reid's Inquiry and Essays*. Indianapolis: Hackett.

Bergin, L. A. (2002). 'Testimony, Epistemic Difference, and Privilege: How Feminist Epistemology Can Improve our Understanding of the Communication of Knowledge', *Social Epistemology*, 16: 197–213.

Bergmann, M. (2004). 'Epistemic Circularity Malignant and Benign', *Philosophy and Phenomenological Research*, 69: 709–27.

—— (2006). *Justification without Awareness*. New York: Oxford University Press.

Bermúdez, J. L. and Millar, A. (eds.) (2002). *Reason and Nature: Essays in the Theory of Rationality*. Oxford: Clarendon Press.

Bigelow, J. and Pargetter, R. (1987). 'Functions', *The Journal of Philosophy*, 84: 181–96.

Bishop, M. A., and Trout, J. D. (2005a). *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford University Press.

—— (2005b). 'The Pathologies of Standard Analytic Epistemology', *Noûs*, 39: 696–714.

Blackburn, S. (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Oxford and New York: Clarendon Press.

Blais, M. (1987). 'Epistemic Tit for Tat', *The Journal of Philosophy*, 82: 335–49.

—— (1988). 'Attitudes and Contents', in Blackburn (1993: 182–97).

—— (1993). *Essays in Quasi-Realism*. Oxford and New York: Oxford University Press.

—— (1996). 'Securing the Nots: Moral Epistemology for the Quasi-Realist', in Sinnott-Armstrong and Timmons (eds.) (1996: 82–100).

—— (1998). *Ruling Passions: A Theory of Practical Reasoning*. Oxford and New York: Clarendon Press.

—— (2006). ‘The Semantics of Non-Factualism’, in Devitt and Hanley (eds.) (2006: 244–52).

Boghossian, P. (2006). *Fear of Knowledge*. Oxford: Oxford University Press.

—— (2008). ‘Replies to Wright, MacFarlane and Sosa’, *Philosophical Studies*, 141: 409–32.

Bohman, J. and Rehg, W. (eds.) (1997). *Deliberative Democracy: Essays on Reason and Politics*. Cambridge, MA: MIT Press.

Bowden, P. (2008). ‘Ecological Thinking, Advocacy and Privilege’, *Ethics, Place & Environment*, 11: 49–90.

Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.

—— (2000a). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

—— (2000b). ‘Fighting Skepticism with Skepticism: Supervaluational Epistemology, Semantic Autonomy, and Natural Kind Skepticism’, *Facta Philosophica*, 2: 163–78.

—— (2009). *Reason in Philosophy: Animating Ideas*. Cambridge, MA: Harvard University Press.

Bratman, M. (1993). ‘Shared Intention’, *Ethics*, 104: 97–113.

Broome, J. (2002). ‘Practical Reasoning’, in Bermúdez and Millar (eds.) (2002: 85–112).

327

Brown, J. (2008). ‘“Subject-Sensitive Invariantism” and the Knowledge Norm of Practical Reasoning’, *Noûs*, 42: 167–89.

Buller, D. (1998). ‘Etiological Theories of Function’, *Biology and Philosophy*, 13: 505–27.

—— (2002). ‘Function and Design Revisited’, in Ariew, Cummins, and Perlman (eds.) (2002: 222–43).

—— (ed.) (1999). *Function, Selection, and Design*. Albany: State University of New York Press.

Burge, T. (1979). ‘Individualism and the Mental’, *Midwest Studies in Philosophy*, 4: 73–121.

—— (1993). ‘Content Preservation’, *The Philosophical Review*, 102: 457–88.

—— (1996). ‘Our Entitlement to Self-Knowledge’, *Proceedings of the Aristotelian Society*, 96: 91–116.



- (1997). ‘Interlocution, Perception, and Memory’, *Philosophical Studies*, 86: 21–47.
- (1999). ‘Comprehension and Interpretation’, in Hahn (ed.) (1999: 229–50).
- (2003). ‘Perceptual Entitlement’, *Philosophy and Phenomenological Research*, 67: 503–48.
- Burling, R. (2005). *The Talking Ape: How Language Evolved*. Oxford: Oxford University Press.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Carruthers, P. and Chamberlain, A. (eds.) (2000). *Evolution and the Human Mind: Language, Modularity and Social Cognition*. New York: Cambridge University Press.
- Laurence, S. and Stich, S. (eds.) (2006). *The Innate Mind: Culture and Cognition*. Oxford: Oxford University Press.
- Stich, S., and Siegal, M. (eds.) (2002). *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.
- Carson, R. (1962). *Silent Spring*. New York & Boston: Houghton Mifflin.
- Carter, J. A. and Chrisman, M. (unpublished). ‘Is Epistemic Expressivism Incompatible with Inquiry?’
- Castoriadis, C. (1991). *Philosophy, Politics, Autonomy: Essays in Political Philosophy*. Ed. David Ames. New York: Oxford University Press.
- (1994). ‘Radical Imagination and the Social Instituting Imaginary’, in Robinson and Rundell (eds.) (1994: 136–54).
- (1998). *The Imaginary Institution of Society*. Trans. Kathleen Blamey. Cambridge, MA: MIT Press.
- Chisholm, R. (1982). *The Foundations of Knowing*. Minneapolis: University of Minnesota Press.
- Chrisman, M. (2007). ‘From Epistemic Contextualism to Epistemic Expressivism’, *Philosophical Studies*, 135: 225–54.
- (2008). ‘Ought to Believe’, *The Journal of Philosophy*, 105: 346–70.
- (2009). ‘Expressivism, Truth, and (Self-)Knowledge’, *Philosophers' Imprint*.
- (forthcoming) ‘Expressivism, Inferentialism and the Theory of Meaning’, in M. Brady (ed.) *New Waves in Meta-Ethics*. Basingstoke: Palgrave Macmillan.
- Christensen, D. (2007). ‘Epistemology of Disagreement: the Good News’, *The Philosophical Review*, 116: 187–217.
- (forthcoming). ‘Higher-Order Evidence’, *Philosophy and Phenomenological Research*.

- and Kornblith, H. (1997). 'Testimony, Memory and the Limits of the *A Priori*', *Philosophical Studies*, 86: 1–20.
- Christiansen, M. and Kirby, S. (eds.) (2003). *Language Evolution*. Oxford: Oxford University Press.
- Clark, A. and Chalmers, D. (1998). 'The Extended Mind', *Analysis*, 58: 7–19.
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Cockburn, C. (1998). *The Space Between Us: Negotiating Gender and National Identities in Conflict*. London: Zed Books.
- Code, L. (1987). *Epistemic Responsibility*. Hanover, NH: University Press of New England.
- (1995). 'Incredulity, Experientialism, and the Politics of Knowledge', in her *Rhetorical Spaces: Essays on (Gendered) Locations*. New York: Routledge, 58–82.
- (2006). *Ecological Thinking: The Politics of Epistemic Location*. New York: Oxford University Press.
- (2008). 'Advocacy, Negotiation and the Politics of Unknowing', *The Southern Journal of Philosophy*, Supplementary Volume: 'Global Feminist Ethics and Politics', 46: 32–51.
- Cohen, J. (1989). 'Deliberation and Democratic Legitimacy', in Hamlin and Pettit (eds.) (1989: 17–34).
- (1997). 'Procedure and Substance in Deliberative Democracy', in Bohman and Rehg (eds.) (1997: 407–37).
- Cohen, S. (1988). 'How to Be a Fallibilist', *Philosophical Perspectives*, 2: 91–123.
- (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Noûs*, 33: 57–89.
- (2002). 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research*, 65: 309–29.
- Coleman, J. and Ferejohn, J. (1986). 'Democracy and Social Choice', *Ethics*, 97: 6–26.
- Conee, E. and Feldman, R. (eds.) (2004). *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.
- Copp, D. (2001). 'Realist-Expressivism: A Neglected Option for Moral Realism', *Social Philosophy and Policy*, 18: 1–43.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Cummins, R. (1975). 'Functional Analysis', *The Journal of Philosophy*, 72, 741–65.

- Cuneo, T. (2008). *The Normative Web*. New York: Oxford University Press.
- Dancy, J. (2004). *Ethics without Principles*. Oxford: Clarendon Press.
- Darwall, S. (2006). *The Second-Person Standpoint*. Cambridge, MA: Harvard University Press.
- Davidson, D. (1980). *Actions and Events*. Oxford: Oxford University Press.
- Davis, W. (2003). *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.
- Delaney, C. (ed.) (1979). *Rationality and Religious Belief*. London: University of Notre Dame Press.
- DeRose, K. (1991). 'Epistemic Possibilities', *The Philosophical Review*, 100: 581–605.
- (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52: 913–29.
- (1995). 'Solving the Skeptical Problem', *The Philosophical Review*, 104: 1–52.
- (1996). 'Knowledge, Assertion and Lotteries', *Australasian Journal of Philosophy*, 74: 568–80.

- DeRose, K. (2002). 'Assertion, Knowledge, and Context', *The Philosophical Review*, 111: 167–203.
- (2009). *The Case for Contextualism: Knowledge, Skepticism and Context*, vol. 1. Oxford and New York: Oxford University Press.
- Descartes, R. (1984). *The Philosophical Writings of Descartes*, vol. 2. Cambridge: Cambridge University Press.
- Dessalles, J.-L. (2007). *Why We Talk: The Evolutionary Origins of Language*. Oxford: Oxford University Press.
- Devitt, M. and Hanley, R. (eds.) (2006). *The Blackwell Guide to the Philosophy of Language*. Oxford: Blackwell.
- Dewey, J. (1927). *The Public and its Problems*. Athens, OH: Swallow Press.
- Dietrich, F. (2007). 'The Premises of Condorcet's Jury Theorem Are Not Simultaneously Justified', *Episteme: A Journal of Social Epistemology*, 5: 56–73.
- and List, C. (2007). 'Judgment Aggregation by Quota Rules: Majority Voting Generalized', *Journal of Theoretical Politics*, 19(4): 391–424.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: Harvard University Press.
- Dummett, M. (1974). *Frege: Philosophy of Language*. Cambridge, MA: Harvard University Press.

—— (1994). ‘Testimony and Memory’, in Matilal and Chakrabarti (eds.) (1994: 251–72).

Dunbar, R. (1996). *Grooming, Gossip, and the Evolution of Language*. Cambridge, MA: Harvard University Press.

Edwards, J. (2003). ‘Burge on Testimony and Memory’, *Analysis*, 60: 124–31.

Elga, A. (2007). ‘Reflection and Disagreement’, *Noûs*, 41: 478–502.

—— (2010). ‘How to Disagree About How to Disagree’, in Feldman and Warfield (eds.) (2010: 175–86).

Elster, J. (1989). *The Cement of Society: A Study of Social Order*. Cambridge: Cambridge University Press.

Estlund, D. (2008). *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.

Fairweather, A. and Zagzebski, L. (eds.) (2001). *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. New York and Oxford: Oxford University Press.

Fallis, D. (2007). ‘Attitudes Toward Epistemic Risk and the Value of Experiments’, *Studia Logica*, 86: 215–46.

—— and Sanger, L. (eds.) (2009). Symposium on ‘The Epistemology of Mass Collaboration’, *Episteme: A Journal of Social Epistemology*, 6: 1–106.

Fantl, J. and McGrath, M. (2002). ‘Evidence, Pragmatics, and Justification’, *The Philosophical Review*, 111: 67–94.

Faulkner, P. (2000). ‘The Social Character of Testimonial Knowledge’, *The Journal of Philosophy*, 97: 581–601.

—— (2002). ‘The Rationality of Our Response to Testimony’, *Synthese*, 131: 353–70.

—— (2007a). ‘On Telling and Trusting’, *Mind*, 116: 875–902.

—— (2007b). ‘A Genealogy of Trust’, *Episteme*, 4: 305–21.

—— (unpublished). ‘Lies and the Problem of Trust’.

Feigl, H. and Scriven, M. (eds.) (1956). *Minnesota Studies in the Philosophy of Science*, vol. 1. Minneapolis, MN: University of Minnesota Press.

Feldman, R. (2000). ‘The Ethics of Belief’, *Philosophy and Phenomenological Research* 60: 667–95. Reprinted in Conee and Feldman (eds.) (2004: 166–96). Page references are to the reprinted version.

—— (2003). *Epistemology*. Upper Saddle River: Prentice Hall, Inc.

—— (2005). ‘Respecting the Evidence’, *Philosophical Perspectives*, 19: 95–119.

- (2006). ‘Epistemological Puzzles About Disagreement’, in Hetherington (ed.) (2006: 216–37).
- (2007). ‘Reasonable Religious Disagreements’, in Antony (ed.) (2007: 194–214).
- (unpublished). ‘Two Epistemic Principles and the Problem of Reasonable Disagreements’.
- and Warfield, T. (eds.) (2010). *Disagreement*. Oxford: Oxford University Press.
- Field, H. (1998). ‘Epistemological Nonfactualism and the a Prioricity of Logic’, *Philosophical Studies*, 92: 1–24.
- (2001). *Truth and the Absence of Fact*. Oxford: Oxford University Press.
- Firth, R. (1978). ‘Are Epistemic Concepts Reducible to Ethical Concepts?’, in Goldman and Kim (eds.) (1978: 215–29).
- Fodor, J. (1990). *A Theory of Content*. Bradford, MA: MIT Press.
- Foley, R. (2006). *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Foucault, M. (1980). *Power/Knowledge: Selected Interviews and Other Writings, 1972–1977*. Ed. Colin Gordon. Trans. Colin Gordon, Leo Marshall, John Mepham, and Kate Soper. New York: Pantheon.
- (1984). ‘Nietzsche, Genealogy, History’, in Rabinow (ed.) (1984: 76–100).
- Fricker, E. (1994). ‘Against Gullibility’, in Matilal and Chakrabarti (1994: 125–61).
- (1995). ‘Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony’, *Mind*, 104: 393–411.
- (2006). ‘Varieties of Anti-Reductionism—A Reply to Goldberg and Henderson’, *Philosophy and Phenomenological Research*, 72, 618–28.
- Fricker, M. (1998). ‘Rational Authority and Social Power: Towards a Truly Social Epistemology’, *Proceedings of the Aristotelian Society*, 98: 159–77.
- (2003). ‘Epistemic Injustice and a Role for Virtue in the Politics of Knowing’, *Metaphilosophy*, 34: 154–73.
- (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- (2008). ‘Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time’, *Philosophical Papers*, 37: 27–50.
- Frye, M. (1996). ‘The Possibility of Feminist Theory’, in Garry and Pearsall (eds.) (1996: 34–47).
- Fuerstein, M. (2009). ‘The Scientific Public: Inquiry in Democratic Society’ (Doctoral dissertation, Columbia University).

Fuller, S. (1988). *Social Epistemology*. Bloomington, IN: Indiana University Press.

Fumerton, R. (1995). *Metaepistemology and Skepticism*. Lanham, MA: Rowman & Littlefield.

—— (2006). ‘The Epistemic Role of Testimony: Internalist and Externalist Perspectives’, in Lackey and Sosa (eds.) (2006: 77–92).

—— (2010). ‘You Can’t Trust a Philosopher’, in Feldman and Warfield (eds.) (2010: 91–110).

331

Fumerton, R. (unpublished). ‘Egocentric Perspectives and the Epistemic Significance of Disagreement’.

Garry, A. and Pearsall, M. (eds.) (1996). *Women, Knowledge, and Reality: Explorations in Feminist Philosophy*. Second edition. New York: Routledge.

Gates, B. T. and Shteir, A. B. (eds.) (1997). *Natural Eloquence: Women Reinscribe Science*. Madison: The University of Wisconsin Press.

Geach, P. T. (1965). ‘Assertion’, *The Philosophical Review*, 74: 449–65.

Gendler, T. S. and Hawthorne, J. (eds.) (2005). *Oxford Studies in Epistemology*, vol. 1. Oxford: Oxford University Press.

—— (forthcoming). *Oxford Studies in Epistemology*, vol. 3. Oxford: Oxford University Press.

Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.

—— (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Gibbons, J. (1996). ‘Externalism and Knowledge of Content’, *Philosophical Review*, 105: 287–310.

Gilbert, D. (1991). ‘How Mental Systems Believe’, *American Psychologist*, 46: 107–19.

—— (1993). ‘The Ascent of Man’, in Wegner and Pennebaker (eds.) (1993: 57–87).

Gilbert, M. (1989). *On Social Facts*. London: Routledge.

Godfrey-Smith, P. (1994). ‘A Modern History Theory of Functions’, *Noûs*, 28: 344–62.

Goldberg, S. (2001). ‘Testimonially Based Knowledge from False Testimony’, *The Philosophical Quarterly*, 51: 512–26.

—— (2006). ‘Testimony as Evidence’, *Philosophica*, 78: 29–52.

—— (ed.) (2007a). *Internalism and Externalism in Semantics and Epistemology*. Oxford: Oxford University Press.

—— (2007b). *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: Cambridge University Press.

—— (2007c). ‘Semantic Externalism and Epistemic Illusions’, in Goldberg (ed.) (2007a: 235–52).

—— (forthcoming a). ‘If that were true I would have heard about it by now’, in Goldman and Whitcomb (eds.) (forthcoming).

—— (forthcoming b). *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.

—— and Henderson, D. (2006). ‘Monitoring and Anti-Reductionism in the Epistemology of Testimony’, *Philosophy and Phenomenological Research*, 72: 576–93.

Goldman, A. I. (1976). ‘Discrimination and Perceptual Knowledge’, *The Journal of Philosophy*, 73: 771–91.

—— (1979). ‘What Is Justified Belief?’, in G. Pappas (1979: 1–24). Reprinted in Goldman (1992b: 105–26).

—— (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

—— (1992a). ‘Epistemic Folkways and Scientific Epistemology’, in Goldman (1992b: 155–77).

—— (1992b). *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge MA: MIT Press.

—— (1994). ‘Argumentation and Social Epistemology’, *The Journal of Philosophy*, 91: 27–49.

—— (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.

332

—— (2001). ‘Experts: Which Ones Should You Trust?’, *Philosophy and Phenomenological Research*, 63: 85–109.

—— (2003). ‘An Epistemological Approach to Argumentation’, *Informal Logic*, 23: 49–61.

—— (2008). ‘Immediate Justification and Process Reliabilism’, in Smith (ed.) (2008: 63–82).

—— (2010). ‘Epistemic Relativism and Reasonable Disagreement’, in Feldman and Warfield (eds.) (2010: 187–215).

—— (forthcoming) ‘Systems-Oriented Social Epistemology’, in Gendler and Hawthorne (eds.) (forthcoming).

—— and Cox, J. (1996). ‘Speech, Truth and the Free Market for Ideas’, *Legal Theory*, 2: 1–32.

- and Kim, J. (eds.) (1978). *Values and Morals*. Dordrecht: Reidel.
- and Olsson, E. (2009). ‘Reliabilism and the Value of Knowledge’, in Haddock, Millar, and Pritchard (eds.) (2009: 19–41).
- and Shaked, M. (1991). ‘An Economic Model of Scientific Activity and Truth Acquisition’, *Philosophical Studies*, 63: 31–55. Reprinted (with proofs of theorems) in Goldman (1992b: 227–55).
- and Whitcomb, D. (eds.) (forthcoming). *The Oxford Handbook of Social Epistemology*. Oxford: Oxford University Press.
- Goldman, R. and Goldman, A. (2009). ‘Review of Larry Laudan's *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*, *Legal Theory*, 15: 55–66.
- Good, I. J. (1967). ‘On the Principle of Total Evidence’, *British Journal for the Philosophy of Science*, 17: 319–21. Reprinted in I. J. Good (1983: 178–80).
- (1983). *The Foundations of Probability and Its Applications*. Minneapolis: University of Minnesota Press.
- Goodman, K. W. (2003). *Ethics and Evidence-Based Medicine: Fallibility and Responsibility in Clinical Science*. Cambridge: Cambridge University Press.
- Graham, P. J. (1997). ‘What is Testimony?’, *The Philosophical Quarterly*, 47: 227–32.
- (1999). ‘Testimony: The Epistemology of Linguistic Acceptance’ (PhD Thesis, Stanford University).
- (2000a). ‘Conveying Information’, *Synthese*, 126: 365–92.
- (2000b). ‘Transferring Knowledge’, *Noûs*, 34: 131–52.
- (2000c). ‘The Reliability of Testimony’, *Philosophy and Phenomenological Research*, 61: 695–708.
- (2006a). ‘Liberal Fundamentalism and Its Rivals’, in Lackey and Sosa (eds.) (2006: 93–115).
- (2006b). ‘Can Testimony Generate Knowledge?’, *Philosophica*, 78: 105–27.
- (2006c). ‘Testimonial Justification: Inferential or Non-Inferential?’, *The Philosophical Quarterly*, 56: 84–95.
- (unpublished). ‘Epistemic Entitlement’.
- Greco, J. (2009). ‘The Value Problem’, in Haddock, Millar, and Pritchard (eds.) (2009: 313–21).
- and Sosa, E. (eds.) (1999). *The Blackwell Guide to Epistemology*. Oxford: Blackwell.
- Grice, H. P. (1957). ‘Meaning’, *The Philosophical Review*, 66: 377–88.
- (1967). ‘Logic and Conversation’, William James Lectures, in Grice (1989).



—— (1989). *Studies in the Way of Words*. Cambridge MA: Harvard University Press.

333

Griffiths, M. and Whitford, M. (eds.) (1988). *Feminist Perspectives in Philosophy*. Bloomington, IN: Indiana University Press.

Guggenheim, D. and Gore, A. (dirs.) (2006). *An Inconvenient Truth: The Planetary Emergency of Global Warming and What We Can Do About It*. Film Presentation. Rodale: Melcher Media.

Gutmann, A. and Thompson, D. (1996). *Democracy and Disagreement*. Cambridge, MA: Harvard University Press.

Gutting, G. (1982). *Religious Belief and Religious Skepticism*. Notre Dame: University of Notre Dame Press.

Habermas, J. (1996). *Between Fact and Norms: Contributions to a Discourse Theory of Law and Democracy*. Trans. W. Rehg. Cambridge, MA: MIT Press.

Haddock, A. and Macpherson, F. (eds.) (2008). *Disjunctivism: Perception, Action, Knowledge*. Oxford: Clarendon Press.

—— Millar, A. and Pritchard, D. H. (eds.) (2009). *Epistemic Value*. Oxford: Oxford University Press.

Hahn, L. E. (ed.) (1999). *The Philosophy of Donald Davidson*. Chicago: Open Court Press.

Hales, S. (2006). *Relativism and the Foundations of Philosophy*. Cambridge, MA: MIT Press.

Hamblin, C. L. (1970). *Fallacies*. London: Methuen.

Hamlin, A. and Pettit, P. (eds.) (1989). *The Good Polity*. Oxford: Blackwell.

Hankinson-Nelson, L. (1990). *Who Knows: From Quine to Feminist Empiricism*. Philadelphia: Temple University Press.

—— (1993). 'Epistemological Communities', in Alcoff and Potter (eds.) (1993: 121–59).

Haraway, D. (1991). 'Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective', in her *Simians, Cyborgs, and Women: The Reinvention of Nature*. New York: Routledge, 183–201.

Harding, S. (1993). 'Rethinking Standpoint Epistemology: What is 'Strong Objectivity'?', in Alcoff and Potter (eds.) (1993: 101–20).

Hardwig, J. (1985). 'Epistemic Dependence', *The Journal of Philosophy*, 82: 335–49.

Hare, R. M. (1952). *The Language of Morals*. Oxford: Clarendon Press.

Harris, P. (2002). 'What do Children Learn from Testimony?', in Carruthers, Stich, and Siegal (eds.) (2002: 316–34).

- Hauser, M. (2000). *Wild Minds*. New York: Henry Holt.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- and Stanley, J. (2008). 'Knowledge and Action', *The Journal of Philosophy*, 105: 571–90.
- Heller, M. (1999). 'The Proper Role for Contextualism in an Anti-Luck Epistemology', *Noûs*, 13: 115–29.
- Hetherington, S. (ed.) (2006). *Epistemology Futures*. Oxford: Oxford University Press.
- Hinchman, E. (2005). 'Telling as Inviting to Trust', *Philosophy and Phenomenological Research*, 70: 562–87.
- (forthcoming). 'The Assurance of Warrant'. *Philosophers' Imprint*.
- (unpublished). 'Assertion, Sincerity, and Knowledge'.
- Hollis, M. and Lukes, S. (eds.) (1982). *Rationality and Relativism*. Cambridge, MA: MIT Press.

- Honderich, T. (ed.) (1985). *Morality and Objectivity: A Tribute to J. L. Mackie*. London: Routledge.
- Hume, D. (1977). *An Enquiry Concerning Human Understanding*, E. Steinberg, ed. Indianapolis: Hackett.
- (2000). *An Enquiry Concerning Human Understanding*, ed. T. L. Beauchamp. Oxford: Clarendon Press.
- Hunter, D. (1998). 'Understanding and Belief', *Philosophy and Phenomenological Research*, 58: 559–80.
- Hurford, J. (2007). *The Origins of Meaning: Language in the Light of Evolution*. Oxford: Oxford University Press.
- Jaswal, V. K. (2004). 'Don't Believe Everything You Hear: Preschoolers' Sensitivity to Speaker Intent in Category Induction', *Child Development*, 75: 1871–85.
- Jeffrey, R. C. (ed) (1980). *Studies in Inductive Logic and Probability*. Berkeley: University of California Press.
- Johnson, R. N. (2003). 'Virtue and Right', *Ethics*, 113: 810–34. Reprinted in *The Philosopher's Annual*, 26.
- Johnston, M. (1989). 'Dispositional Theories of Value', *The Aristotelian Society*, Supplementary Volume, 63: 139–74.
- Jordan, J. and Howard-Snyder, D. (eds.) (1996). *Faith, Freedom, and Rationality: Philosophy of Religion Today*. London: Rowman and Littlefield.

Kaplan, M. (2003). 'Critical Study: Who Cares What You Know?', *The Philosophical Quarterly*, 53: 105–16.

Kelly, T. (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research*, 66: 612–40.

—— (2005). 'The Epistemic Significance of Disagreement', in Gendler and Hawthorne (eds.) (2005: 167–97).

—— (2010). 'Peer Disagreement and Higher Order Evidence', in Feldman and Warfield (eds.) (2010: 111–74).

Koenig, M., Clement, F., and Harris, P. (2004). 'Trust in Testimony: Children's Use of True and False Statements', *Psychological Science*, 10: 694–8.

—— and Harris, P. (2005). 'The Role of Social Cognition in Early Trust', *Trends in Cognitive Science*, 9: 457–9.

Kolbel, M. (2003). 'Faultless Disagreement', *Proceedings of the Aristotelian Society*, 104: 53–73.

Kornblith, H. (ed) (1994a). *Naturalizing Epistemology*, 2nd edn. Cambridge, MA: MIT Press.

—— (1994b). 'A Conservative Approach to Social Epistemology', in Schmitt (ed.) (1994: 93–110).

—— (2001). 'Epistemic Obligation and the Possibility of Internalism', in Fairweather and Zagzebski (eds.) (2001: 231–48).

—— (2002). *Knowledge and its Place in Nature*. Oxford: Clarendon Press.

—— (2008). 'Knowledge Needs No Justification', in Smith (ed.) (2008: 5–23).

Kornhauser, L. and Sager, L. G. (1993). 'The One and the Many: Adjudication in Collegial Courts', *California Law Review*, 81: 1–59.

Kroll, G. (2002). 'Ecology as a Subversive Subject', *Reflections: Newsletter of the Program for Ethics, Science and the Environment*, Oregon State University, 9: 17–21.

335

Kusch, M. (2002). *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Clarendon Press.

Kusch, M. (2009). 'Testimony and the Value of Knowledge', in Haddock, Millar, and Pritchard (eds.) (2009: 60–94).

Kvanvig, J. (1998). 'Why Should Inquiring Minds Want to Know?', *The Monist*, 81: 426–51.

—— (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

- Lackey, J. (2005). 'Memory as a Generative Source', *Philosophy and Phenomenological Research*, 70: 636–58.
- (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese*, 158: 345–61.
- (2008). *Learning from Words*. Oxford: Oxford University Press.
- (forthcoming). 'What Should We Do When We Disagree', in Gendler and Hawthorne (eds.) (forthcoming).
- and Sosa, E. (eds.) (2006). *The Epistemology of Testimony*. Oxford: Clarendon Press.
- Lane, M. (1999). 'States of Nature, Epistemic and Political', *Proceedings of the Aristotelian Society*, 99: 211–24.
- Lasersohn, P. (2005). 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy*, 28: 643–86.
- Latour, B. (1987). *Science in Action*. Cambridge, MA: Harvard University Press.
- (2004). *The Politics of Nature: How to Bring the Sciences into Democracy*, translated by Catherine Porter. Cambridge, MA: Harvard University Press.
- and Woolgar, S. (1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press.
- Laudan, L. (2006). *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. New York: Cambridge University Press.
- Lear, L. J. (1977). *Rachel Carson: Witness for Nature*. London: Penguin.
- Lehrer, K. (1990). *Theory of Knowledge*. Boulder: Westview.
- (1997). *Self-Trust: A Study of Reason, Knowledge, and Autonomy*. Oxford: Oxford University Press.
- and Wagner, C. (1981). *Rational Consensus in Science and Society*. Dordrecht: Reidel.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Lewis, D. (1969). *Convention – A Philosophical Study*. Cambridge, MA: Harvard University Press.
- (1980). 'A Subjectivist's Guide to Objective Chance', in Jeffrey (ed.) (1980: 263–93).
- (1989). 'Dispositional Theories of Value', *The Aristotelian Society*, Supplementary Volume, 63: 113–37.
- (1994). 'Humean Supervenience Debugged', *Mind*, 103: 473–90.
- (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549–67.

Lipton, P. (1998). 'The Epistemology of Testimony', *Studies in the History and Philosophy of Science*, 29: 1–31.

—— (2007). 'Alien Abduction', *Episteme*, 4: 238–51.

List, C. (forthcoming). 'Group Communication and the Transformation of Judgments: An Impossibility Result', *The Journal of Political Philosophy*.

336

—— and Pettit, P. (2002). 'Aggregating Sets of Judgments: An Impossibility Result', *Economics and Philosophy*, 18: 89–110.

—— — (forthcoming). *Group Agency, The Possibility, Design and Status of Corporate Agents*.

Locke, J. (1975). *An Essay Concerning Human Understanding*, ed. P. H. Nidditch. Oxford: Clarendon Press.

Lorenzen, P. (2001). 'Basic Objectives of Dialogue Logic in Historical Perspective', *Synthese*, 127: 255–63.

Ludlow, P. (1995). 'Social Externalism, Self-Knowledge, and Memory', *Analysis*, 55: 232–7.

Lynch, M. P. (2004). *True to Life*. Cambridge, MA: MIT Press.

—— (2009a). Truth, Value and Epistemic Expressivism. *Philosophy and Phenomenological Research*, 79: 76–97.

—— (2009b). 'The Truth of Values and the Values of Truth', in Haddock, Millar, and Pritchard (eds.) (2009: 225–42).

Lytle, M. H. (2007). *The Gentle Subversive: Rachel Carson, Silent Spring, and the Rise of the Environmental Movement*. New York: Oxford University Press.

McDowell, J. (1980). 'Meaning, Communication, and Knowledge', in van Straaten (ed.) (1980: 117–39).

—— (1981). 'Anti-Realism and the Epistemology of Understanding', in Parret and Bouveresse (eds.) (1981). Reprinted in McDowell (1998a: 314–43).

—— (1982). 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, 68: 455–79. Reprinted in McDowell (1998a: 369–94).

—— (1984). 'Wittgenstein on Following a Rule', *Synthese*, 58: 325–63. Reprinted in McDowell (1998b: 221–62).

—— (1985). 'Values and Secondary Qualities', in Honderich (ed.) (1985: 110–29).

—— (1998a). *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

—— (1998b). *Mind, Value, Reality*. Cambridge, MA: Harvard University Press.

- MacFarlane, J. (2005). 'The Assessment Sensitivity of Knowledge Attributions', in Gendler and Hawthorne (eds.) (2005: 197–233).
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. London: Penguin.
- McNaughton, D. (1988). *Moral Vision*. Oxford: Blackwell.
- Maher, P. (1990). 'Why Scientists Gather Evidence', *The British Journal for the Philosophy of Science*, 41: 103–19.
- Malmgren, A-S. (2006). 'Is There A *Priori* Knowledge by Testimony?', *The Philosophical Review*, 115: 199–241.
- Matilal, B. K. and Chakrabarti, A. (eds.) (1994). *Knowing from Words*. Amsterdam: Kluwer Academic Publishers.
- Messing, K. (1998). *One-Eyed Science: Occupational Health and Women Workers*. Philadelphia: Temple University Press.
- Mill, J. S. (1962). 'On Liberty', in Warnock (ed.), *John Stuart Mill: Utilitarianism, On Liberty, Essay on Bentham*. New York: Meridian.
- Millar, A. (2000). 'The Scope of Perceptual Knowledge', *Philosophy*, 75: 73–88.
- (2004). *Understanding People: Normativity and Rationalizing Explanation*. Oxford: Clarendon Press.

- Millar, A. (2007a). 'What the Disjunctivist is Right About', *Philosophy and Phenomenological Research*, 74: 176–98.
- (2007b). 'The State of Knowing', in E. Sosa and E. Villanueva (eds.) *Philosophical Issues, 17, The Metaphysics of Epistemology*: 179–96.
- (2008). 'Perceptual-Recognitional Abilities and Perceptual Knowledge', in Haddock and Macpherson (eds.) (2008: 330–47).
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- (2002). 'Biofunctions: Two Paradigms', in Ariew, Cummins, and Perlman (2002: 113–43).
- (2004). *Varieties of Meaning*. Cambridge, MA: MIT Press.
- (2005). *Language: A Biological View*. New York: Oxford University Press.
- Mills, C. (2008). 'White Ignorance', in Proctor and Schiebinger (eds.) (2008: 230–49).
- Misak, C. (2000). *Truth, Politics, Morality: Pragmatism and Deliberation*. London: Routledge.
- Moffett, M. (2007). 'Reasonable Disagreement and Rational Group Inquiry'. *Episteme*, 4: 352–67.

Moore, C., Bryant, D., and Furrow, D. (1989). 'Mental Terms and the Development of Certainty', *Child Development*, 60: 167–71.

Moore, G. E. (1939). 'Proof of an External World', *Proceedings of the British Academy*, 25: 273–300. Reprinted in Moore (1959b: 127–50) to which page references refer.

—— (1959a). 'Certainty', Howison lecture delivered at the University of California, Berkeley, in 1941. Printed in his (1959b: 227–51).

—— (1959b). *Philosophical Papers*. London: George Allen & Unwin.

Moran, R. (2005). 'Problems of Sincerity', *Proceedings of the Aristotelian Society*. 105: 341–61.

—— (2006). 'Getting Told and Being Believed', in Lackey and Sosa (eds.) (2006: 272–306).

—— (unpublished). 'Testimony and the Second-Person'.

Mulhall, S. (1990). *On Being in the World: Wittgenstein and Heidegger on Seeing Aspects*. London: Routledge.

Nagel, T. (1997). *The Last Word*. Oxford: Oxford University Press.

Neander, K. (1991a). 'Functions as Selected Effects', *Philosophy of Science*, 58: 168–84.

—— (1991b). 'The Teleological Notion of 'Function'', *Australasian Journal of Philosophy*, 69: 454–68.

Neta, R. (2003). 'Contextualism and the Problem of the External World', *Philosophy and Phenomenological Research*, 66: 1–31.

—— (2006). 'Epistemology Factualized: New Contractarian Foundations for Epistemology', *Synthese*, 150: 247–80.

—— (2008). 'Anti-Intellectualism and the Knowledge-Action Principle', *Philosophy and Phenomenological Research*, 75: 180–7.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Nuccetelli, S. and Seay, G. (eds.) (2007). *Themes from G. E. Moore*. Oxford: Oxford University Press.

Origg, G. and Sperber, D. (2000). 'Evolution, Communication and the Proper Function of Language', in Carruthers and Chamberlain (eds.) (2000: 140–69).

O'Rourke, M., Campbell, J., and Silverstein, H. (eds.) (2005). *Contemporary Topics in Philosophy, Vol. 5: Knowledge and Skepticism*. Cambridge, MA: MIT Press.

Owens, D. (2000). *Reason without Freedom: The Problem of Epistemic Normativity*. London: Routledge.

- (2006). ‘Testimony and Assertion’, *Philosophical Studies*, 130: 105–29.
- Pappas, G. (ed.) (1979). *Justification and Knowledge*. Dordrecht: Reidel.
- Parret, H. and Bouveresse, J. (eds.) (1981). *Meaning and Understanding*. Berlin: Walter de Gruyter.
- Pettit, P. (1990). ‘Virtus Normativa: Rational Choice Perspectives’, *Ethics*, 100: 725–55.
- (2003). ‘Groups with Minds of Their Own’, in Schmitt (ed.) (2003: 172–5).
- Pinker, S. (2003). ‘Language as an Adaptation to a Cognitive Niche’, in Christiansen and Kirby (eds.) (2003: 16–37).
- and Bloom, P. (1990). ‘Natural Language and Natural Selection’, *Behavioral and Brain Sciences*, 13: 707–84.
- Plantinga, A. (1993). *Warrant and Proper Function*. New York: Oxford University Press.
- Pollock, J. and Cruz, J. (1999). *Contemporary Theories of Knowledge*. Lanham, MD, Rowman & Littlefield.
- Popper, K. (1962). *Conjectures and Refutations: The Growth of Scientific Knowledge*. New York: Basic Books.
- Preyer, G. and Peter, G. (ed.) (2005). *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Price, H. H. (1969). *Belief*. London: George Allen and Unwin.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- (2007a). ‘The Value of Knowledge’. Available at: <http://www.philosophy.stir.ac.uk/postgraduate/documents/ValueOfKnowledge2.pdf>.
- (2007b). ‘Recent Work on Epistemic Value’, *American Philosophical Quarterly*, 44: 85–110.
- Millar, A., and Haddock A. (2010). *The Nature and Value of Knowledge*. Oxford: Oxford University Press.
- Proctor, R. N. and Schiebinger, L. (eds.) (2008). *Agnotology: The Making and Unmaking of Ignorance*. Stanford: Stanford University Press.
- Quine, W. V. O. (1969a). *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- (1969b). ‘Epistemology Naturalized’, in Quine (1969a: 69–90).
- (1994). ‘Natural Kinds’, in Kornblith (ed.) (1994a: 15–32).
- Rabinow, P. (ed.) (1984). *The Foucault Reader: An Introduction To Foucault's Thought*. London: Penguin.



Raglon, R. (1997). 'Rachel Carson and her Legacy', in Gates and Shteir (eds.) (1997: 196–211).

Rawls, J. (1993). *Political Liberalism*. New York: Columbia University Press.

— (1997). 'The Idea of Public Reason Revisited', *University of Chicago Law Review*, 64: 765–807.

Reed, B. (2006). 'Epistemic Circularity Squared? Skepticism about Common Sense', *Philosophy and Phenomenological Research*, 73: 186–97.

Reid, T. (1983). *Essay on the Intellectual Powers of Man*, in Beanblossom and Lehrer (eds.) (1983).

339

Ridge, M. (2006). 'Ecumenical Expressivism: Finessing Frege', *Ethics*, 16: 302–36.

— (2009). 'Truth in Ecumenical Expressivism', in Sobel and Wall (eds.) (2009: 219–42).

Robinson, G. and Rundell, J. (eds.) (1994). *Rethinking Imagination: Culture and Creativity*. London: Routledge.

Rorty, R. (1979). *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

— (1995). 'Is Truth a Goal of Enquiry? Davidson Vs. Wright', *Philosophical Quarterly*, 45: 281–300.

Rosen, G. (2001). 'Nominalism, Naturalism, Epistemic Relativism', *Philosophical Perspectives*, 15: 60–91.

Rosenberg, J. (1974). *Linguistic Representation*. Dordrecht: D. Reidel.

— (2002). *Thinking about Knowing*. New York and Oxford: Clarendon Press and Oxford University Press.

Ross, A. (1986). 'Why Do We Believe What We are Told?', *Ratio*, 28: 69–88.

Rysiew, P. (2001). 'The Context-Sensitivity of Knowledge Attributions', *Noûs*, 35: 477–514.

— (2007). 'Speaking of Knowing', *Noûs*, 41: 627–62.

Sackett, D. L., Straus, S. E., Richardson, W. S., Rosenberg, W., and Haynes, R. B. (2000). *Evidence-Based Medicine: How to Practice and Teach EBM*, 2nd edn. Edinburgh, London, New York: Churchill Livingstone.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

Schiffer, S. (1972). *Meaning*. Oxford: Oxford University Press.

Schmitt, F. (1994a). 'The Justification of Group Beliefs', in Schmitt (1994b: 257–87).

- (ed.) (1994b). *Socializing Epistemology*. Lanham, Maryland: Rowman & Littlefield.
- (1999). ‘Social Epistemology’, in Greco and Sosa (eds.) (1999: 354–82).
- (ed.) (2003). *Socializing Metaphysics*. Lanham, MD: Rowman & Littlefield.
- (2006). ‘Testimonial Justification and Transindividual Reasons’, in Lackey and Sosa (eds.) (2006: 193–224).
- Schroeder, M. (2008). *Being For*. Oxford and New York: Oxford University Press.
- Schwartz, P. (2002). ‘The Continuing Usefulness Account of Proper Function’, in Ariew, Cummins, and Perlman (eds.) (2002: 244–60).
- Sellars, W. (1956). ‘Empiricism and the Philosophy of Mind’, in Feigl and Scriven (1956: 253–329).
- Seller, A. (1988). ‘Realism versus Relativism: Towards a Politically Adequate Epistemology’, in Griffiths and Whitford (eds.) (1988: 169–86).
- Shafer-Landau, R. (ed.) (2009). *Oxford Studies in Metaethics*, vol. 4. Oxford and New York: Oxford University Press.
- Shapin, S. (1994). *A Social History of Truth*. Chicago: The University of Chicago Press.
- and Schaffer, S. (1985). *Leviathan and the Air-Pump: Hobbes, Boyle and the Experimental Life*. Princeton: Princeton University Press.
- Shope, R. (1978). ‘The Conditional Fallacy in Contemporary Philosophy’, *The Journal of Philosophy*, 75: 397–413.
- Sideris, L. H. and Moore, K. D. (eds.) (2008). *Rachel Carson: Legacy and Challenge*. Albany: State University of New York Press.

- Sinnott-Armstrong, W. and Timmons, M. (eds.) (1996). *Moral Knowledge: New Readings in Moral Epistemology*. Oxford and New York: Oxford University Press.
- Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.
- Smith, M. (1987). ‘The Humean Theory of Motivation’, *Mind*, 96: 36–61.
- (1994). *The Moral Problem*. Oxford: Blackwell.
- Smith, M. B. (2001). ‘“Silence, Miss Carson!”: Science, Gender, and the Reception of *Silent Spring*’, *Feminist Studies*, 27: 733–52. Reprinted in Sideris and Moore (eds.) (2008: 168–89).
- Smith, Q. (ed.) (2008). *Epistemology: New Essays*. Oxford: Oxford University Press.
- Sobel, D. and Wall, S. (ed.) (2009). *Reasons for Action*. Cambridge: Cambridge University Press.

Sober, E. (1994a). 'The Primacy of Truth-Telling and the Evolution of Lying', in Sober (1994b: 71–92).

—— (1994b). *From a Biological Point of View*. Cambridge: Cambridge University Press.

Sosa, E. (1980). 'The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge', *Midwest Studies in Philosophy*, 5: 3–25.

—— (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.

—— (1997a). 'How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes', *Philosophical Studies*, 85: 229–49.

—— (1997b). 'Reflective Knowledge in the Best of Circles', *The Journal of Philosophy*, 94: 410–30.

—— (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. I, Oxford and New York: Oxford University Press.

—— (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Vol. II, Oxford and New York: Oxford University Press.

Sperber, D. (2000a). 'Metarepresentations in an Evolutionary Perspective', in Sperber (ed.) (2000b: 117–37).

—— (ed.) (2000b). *Metarepresentations: A Multidisciplinary Perspective*. New York: Oxford University Press.

—— (2001). 'An Evolutionary Perspective on Testimony and Argumentation', *Philosophical Topics*, 29: 401–13.

Sripada, C. S. and Stich, S. (2006). 'A Framework for the Psychology of Norms', in Carruthers, Laurence and Stich (eds.) (2006: 280–301).

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford and New York: Clarendon Press.

Steup, M. (2004). 'Internalist Reliabilism', *Philosophical Issues*, 14: 403–25.

Stevenson, C. L. (1937). 'Emotive Meaning of Moral Terms', *Mind*, 46: 14–31.

Stich, S. (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.

Stroud, B. (1994). 'Skepticism, "Externalism" and the Goal of Epistemology', *The Aristotelian Society*, Supplementary Volume, 68: 291–307.

Stump, E. and Murray, M. (eds.) (1999). *Philosophy of Religion: The Big Questions*. Malden, MA: Wiley-Blackwell.

Sullivan, S. and Tuana, N. (eds.) (2006). *Hypatia: A Journal of Feminist Philosophy* Special Issue: Feminist Epistemologies of Ignorance.

- Sullivan, S. and Tuana, N. (eds.) (2007). *Race and Epistemologies of Ignorance*. Albany: State University of New York Press.
- Talisse, R. (2005). *Democracy After Liberalism: Pragmatism and Deliberative Politics*. London: Routledge.
- Tehan v. United States* (1966) 383 U.S. 406, 416.
- Timmons, M. (1999). *Morality without Foundations: A Defense of Ethical Contextualism*. New York: Oxford University Press.
- Tomasello, M. (2003). *Constructing a Language*. Cambridge, MA: Harvard University Press.
- (2008). *Origins of Human Communication*. Cambridge, MA: MIT Press.
- Toulmin, S. (1958). *The Uses of Argument*. Cambridge: Cambridge University Press.
- Travis, C. (2005). ‘A Sense of Occasion’, *The Philosophical Quarterly*, 55: 286–314.
- Tuomela, R. (1990). ‘Can Collectives Have Beliefs?’, *Acta Philosophica Fennica*. 49, Special Issue: Language, Knowledge, and Intentionality: Perspectives on the Philosophy of Jaakko Hintikka, ed. Haaparanta, Kusch, and Niiluoto 454–72.
- Unger, P. (1968). ‘An Analysis of Factual Knowledge’, *The Journal of Philosophy*, 65: 157–70.
- (1975). *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- (1984). *Philosophical Relativity*. Minneapolis: University of Minnesota Press.
- Van Inwagen, P. (1996). ‘It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything on Insufficient Evidence’, in Jordan and Howard-Snyder (eds.) (1996: 137–153). Reprinted in Stump and Murray (eds.) (1999: 273–85).
- (2010). ‘We’re Right, They’re Wrong’ in Feldman and Warfield (eds.) (2010: 10–28).
- Van Straaten, Z. (ed.) (1980). *Philosophical Subjects*. Oxford: Clarendon Press.
- Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wegner, D. M. and Pennebaker, J. W. (eds.) (1993). *Handbook of Mental Control*. Englewood Cliffs, NJ: Prentice-Hall.
- Weinberg, J., Nichols, S., and Stich, S. (2001). ‘Normativity and Epistemic Intuitions’, *Philosophical Topics*, 29: 429–60.
- Welbourne, M. (1979). ‘The Transmission of Knowledge’, *The Philosophical Quarterly*, 29: 1–9.
- (1986). *The Community of Knowledge*. Aberdeen: Aberdeen University Press.

—— (2001). *Knowledge*. Chesham, Bucks: Acumen.

White, R. (2005). 'Epistemic Permissiveness', *Philosophical Perspectives*, 19: 445–59.

Wiggins, D. (1987). *Needs, Values, Truth*. Oxford: Oxford University Press.

Williams, B. (1973). 'Deciding to Believe', in his *Problems of the Self: Philosophical Papers 1956–1972*. Cambridge: Cambridge University Press, 136–51.

—— (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.

Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.

—— (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford University Press.

—— (2008). 'Responsibility and Reliability', *Philosophical Papers*, 37: 1–26.

342

Williams, P. J. (1991). *The Alchemy of Race and Rights: Diary of a Law Professor*. Cambridge, MA: Harvard University Press.

—— (1997). *Seeing a Color-Blind Future: The Paradox of Race*. New York: Noonday Press.

Williamson, T. (1996). 'Knowing and Asserting', *The Philosophical Review*, 105: 489–523.

—— (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

—— (2005). 'Knowledge, Context, and the Agent's Point of View', in Preyer and Peter (eds.) (2005: 91–114).

Wittgenstein, L. (1969). *On Certainty*. Ed. G. E. M. Anscombe and G. H. von Wright. Trans. D. Paul and G. E. M. Anscombe. Oxford: Basil Blackwell.

Woods, J., Irvine, A., and Walton, D. (2000). *Argument: Critical Thinking, Logic and the Fallacies*. Toronto: Prentice-Hall.

Wright, C. (2004). 'Warrant for Nothing (and Foundations for Free)?', *The Aristotelian Society*, Supplementary Volume, 78: 167–212.

—— (2007). 'The Perils of Dogmatism', in Nuccetelli and Seay (eds.) (2007: 25–48).

—— (2009). 'Fear of Relativism', *Philosophical Studies*, 141: 379–90.

Wright, L. (1973). 'Functions', *The Philosophical Review*, 82: 139–68.

Zagzebski, L. (1996) *Virtues of the Mind*. Cambridge: Cambridge University Press.

343

344

## Index

accuracy 96–7 , 140–6

Adler, J. 131n. , 135n. , 147n. , 192n.

advocacy 45–50

aggregation function/procedure 16–17

agnotology 48

Almassi, B 314n.

Alston, W. 1–2 , 4 , 12 , 14n. , 18 , 25–6 , 148 , 151n. , 263n. , 267

Anderson, E. 24–5 , 36

Annis, D. 12

Anscombe, G. E. M. 130n.

argument from highly reliable effects 167

argument from occasion specific intentions 163 , 166n.

argument from social interaction 168

argument from successful communication 164–5 , 166n.

assertion 11–13 , 90–1 , 113n. , 126n. , 148–74 , 179n. , 222–3 , 231–42

Audi, R. 148 , 153

Austin, J. L. 127n. , 130n. , 188n. , 223n. , 288–9

Axelrod, R. 136n. , 138

Axtell, G. 54n.

Ayer, A. J. 114n. , 115n. , 118–21

Bach, K. 151n. , 164

Barnes, B. 3–4

Bar-On, D. 114n.

Barth, E. 13n.

Bayesian approach 7–8

belief(s)

degree of 9

reasonable 9–10 , 292

socially acquired 171 *see also* comprehension (based beliefs), perceptual beliefs

Bergin, L. 32n.  
Bergmann, M. 150n. , 263n.  
Bigelow, J. 155n.  
Bishop, M. 70n.  
Blackburn, S. 116n. , 119n.  
Blais, M. 136n.  
Bloom, P. 168  
Bloor, D. 3–4  
Boghossian, P. 5  
Brandom, R. 12n. , 54–6 , 59n. , 63n. , 113n. , 116 , 120n. , 121n. , 128 , 223n.  
Bratman, M. 15  
Broome, J. 180n.  
Brown, J. 126n.  
Bryant, D. 152  
Buller, D. 154  
Burge, T. 149–50 , 182n. , 188n. , 194–215 , 246 , 248–9 , 255n.  
Burling, R. 168–71 , 173  
Carnap, R. 114n.  
Carson, R. 38–40 , 46–9  
Castoriadis, C. 35  
Chalmers, D 258  
charity, principle of 216  
Chisholm, R. 199  
Chrisman, M. 114n. , 116 , 118n. , 125n.  
Christensen, D. 7 , 11 , 25n. , 27 , 194n. , 255n. , 279n. , 282n. , 301–2 , 307n. , 309n. , 313n. , 314n. , 315–16 , 321 , 323–4  
Clark, A. 258  
Clement, F. 152  
closure 17  
Coady, C. A. J. 7 , 149 , 160–3 , 171  
Code, L. 29n. , 31n. , 35n. , 38n.

Cockburn, C. 46n.  
cognitive conflict 172  
cognitive equality 302–6 , 311–12  
cognitive science 120–2  
Cohen, S. 24 , 263n.  
Coleman, J. 23  
communication 11–12 , 32 , 148–74 , 175–6 , 182 , 188n. , 220  
competence 23 , 189–90 , 231n. , 282–3 , 285n. , 293  
comprehension 148–74 , 248  
based beliefs 149–51 , 173–4  
with filtering 150–3 , 159–62 , 170–4  
commitments 47 , 179–80 , 221  
conceptual synthesis 81–5  
confirmation 7 , 20  
content preservation 246 , 253–4  
contextualism 5 , 12 , 59 , 63–5 , 91 , 116–17  
proto 64–65  
conventions 82 , 133–7  
cooperative outcomes 135–8  
cooperative principle 131–4 , 140–2 , 150  
coordination equilibria 133–8  
coordination problem 133–4

345

Copp, D. 118n.  
Cox, J. 23n.  
Craig, E. 12 , 30 , 54 , 59–62 , 67 , 69–88 , 90n. , 116 , 127  
criterion argument 262–3  
Cruz, J. 113n. , 126n.  
Cummins, R. 154n.  
Cuneo, T. 118n.  
Dancy, J. 93n.



Darwall, S. 224n.  
Davidson, D. 115n.  
Davis, W. 114n.  
decision theory 7–8 , 27 , 97  
deconstructionism 3  
defeater(s) 13 , 19 , 91 , 233 , 310 , 320  
democracy 23–7  
DeRose, K. 10 , 12n. , 17 , 113n. , 223n.  
Descartes, R. 288–9  
Dessalles, J. L. 168–73  
Dewey, J. 24  
Dietrich, F. 17 , 23n.  
disagreement 6–7 , 11 , 14 , 263–8 , 270–7 , 298–325  
conformist view of 300 , 315–18  
idealized 303–15  
justificationist view of 318–25  
nonconformist view of 299–314  
ordinary 304–19 , 321  
reasonable 268n. , 278–81 , 292–5 , 299 , 305 , 312–13 *see also* epistemic  
incommensurability  
disclosure 279–80 , 292–5 , 300–7 , 311–16  
doxastic decision-makers 2 , 6 , 11 , 15–16 , 19–20  
doxastic revision 299–305 , 309–14 , 319–24  
Dretske, F. 120n.  
Dummett, M. 121n. , 125 , 255n.  
Dunbar, R. 168–73  
ecological thinking 35–45  
Edwards, J. 255n.  
Elga, A. 7 , 282n. , 284n. , 299n. , 300n. , 301–4 , 309n. , 313 , 322–4  
Elster, J. 129 , 139  
enquirer 60–4 , 86

enquiry 7 , 30 , 74–80 , 178n.  
method of 75  
stopper 74–9 *see also* norms (of enquiry)  
entitlement(s) 54 , 64 , 188n. , 194–203 , 208 , 212–15 , 218n. , 221n. , 234–5 , 271  
a priori 194 , 199 , 203 , 206–8 , 212–13  
externalism about 196  
preservation of 194 , 199–203 , 212  
prima facie 149–50 , 159 , 170  
pro tanto 149–50 , 159 , 170  
testimonial 148–74 *see also* justification, warrant  
epistemic agents 2 , 3 , 10  
epistemic angst 273  
epistemic autonomy 45  
epistemic circularity 262–74  
epistemic community 30 , 51 , 59–60  
epistemic consequentialism 22  
epistemic division of labour 51–2 , 91  
epistemic incommensurability 262–77 .  
epistemic inferentialism 124–7  
epistemic interdependence 51  
epistemic justice / injustice 33–4 , 37 , 41–5 , 49–50 , 53  
epistemic method game 275–6  
epistemic nihilism 5  
epistemic outcomes 15 , 18  
epistemic peer 282–5 , 299–325  
epistemic practices 37–8 , 47 , 57  
epistemic principles 7 , 264–75  
epistemic reliance 245–9 , 256  
epistemic responsibility 56–7 , 63 , 219–20 , 228–33 , 237–40  
epistemic rules 55  
epistemic self-reliance 29–31

epistemic State of Nature 59–67 , 69–71 , 76–7 , 81–7 , 140–1 , 144–7  
epistemic subject  
abstracted social conception of 51–3 , 57  
situated social conception of 52–3  
epistemic utility 8n. , 181  
epistemic value 98 , 116 , 221n. , 226n. *see also* knowledge (value of), value  
epistemological context 67–8  
epistemological realism 58–9  
epistemology  
applied 25  
communitarian  
feminist 36  
meta 92 , 122–3 , 127–8  
methodology of  
naturalized 35–7 , 259–60  
of silence 243–61  
Anglo-American/real/mainstream/standard/traditional 1–7 , 15–18 , 23–6 , 29–36 , 39 ,  
43 , 51 , 70  
virtue 33 , 37 , 113n. *see also* social epistemology.  
equal weight view 283–5 , 301 , 313–15 , 321–2  
Estlund, D. 24

346

etioloical functions 106–9 , 150n. , 153–74  
and normal conditions 150–9  
and normal functioning 150–9 , 169n.  
evidence 33–44 , 176–7 , 183–6 , 192 , 278–97 , 298–315  
social 6–7 , 17–19  
evidentialism 9–10 , 148 , 216–19  
evidential / nonevidential epistemic reason 222–7 , 231–7  
evidential equality 302–6 , 315  
expert(s) *see* layman vs. expert

explanation(s) 71–2 , 81–5 , 90–5 , 155n. , 269 , 271–3 , 324  
assurance 146  
attitude based 138–40  
behaviourial 138–9 *see also* practical explications  
experiment/experimentation 38–9  
expressivism 112–28  
ethical 113–22  
epistemic 116–26  
and queerness 117 , 124–5 *see also* Frege-Geach problem  
extended mind thesis 256–60  
Fantl, J. 10  
Fallis, D. 8n. , 19n.  
Faulkner, P. 135n. , 140n. , 144n. , 149 , 160n. , 165  
Feldman, R. 7 , 9–10 , 279n. , 282n. , 284n. , 286n. , 300–5 , 307n. , 311 , 316 , 321n.  
feminism/feminist theory *see* epistemology (feminist)  
Ferejohn, J. 23  
Field, H. 116n. , 272n.  
filtering 149–53 , 159–74 *see also* comprehension (with filtering)  
Firth, R. 218n.  
Fodor, J. 120n.  
Foley, R. 257n.  
Foucault, M. 35 , 48n. , 66–7  
freedom of speech 22 , 26  
Frege-Geach problem 118–22  
Fricker, E. 7 , 149 , 188n. , 189n.  
Fricker, M. 30 , 33–5 , 40–2 , 53n. , 54n. , 70n. , 184n. , 187n.  
Fuerstein, M. 25n.  
Fuller, S. 3n.  
Fumerton, R. 177n. , 184n. , 263n. , 300n.  
Furrow, D. 152  
Frye, M. 47

game theory 27 , 139–40 , 143–5  
Geach, P. T. 118  
*Geisteswissenschaften* 37  
*Gestalt* 183  
Gettier cases 149n. , 203  
Gibbons, J. 253n.  
Gibbard, A. 114–16 , 119n.  
Gilbert, D. 152n.  
Gilbert, M. 15  
Godfrey-Smith, P. 154n. , 157  
Goldberg, S. 151n. , 153 , 188n. , 216n. , 245n. , 247n. , 249n. , 250n. , 251n. , 253n. ,  
255n. , 256n.  
Goldman, A. I. 1 , 4n. , 8n. , 11 , 14 , 21–3 , 27n. , 31n. , 70n. , 99–103 , 148 , 170 , 186n.  
, 189 , 246n. , 311n.  
Goldman, R. 27n.  
Good, I. J. 8  
Goodman, K. W. 43n.  
Gore, A. 48  
Graham, P. J. 149n. , 151n. , 152n. , 159n. , 161  
Greco, J. 92  
Grice, H. P. 12 , 114 , 131–2 , 142  
Guggenheim, D. 48n.  
Gutmann, A. 24  
Gutting, G. 282n. , 299n. , 302n.  
Habermas, J. 24  
Haddock, A. 116n. , 178n. , 183n.  
Hales, S. 272n.  
Hamblin, C. 13n.  
Hankinson-Nelson, L. 3 , 51n.  
Haraway, D. 35n.  
Harding, S. 38n.

Hardwig, J. 51n. , 52  
Hare, R. M. 119n.  
Harnish, R. 151n.  
Harris, P. 152  
Hauser, M. 172n.  
Hawthorne, J. 12n. , 90n. , 113n. , 126n.  
Heller, M. 116n.  
Henderson, D. 153 , 188n. , 255n.  
Herrmann, E. 275n.  
Hinchman, E. 217n. , 218n. , 221n. , 222–3 , 233n. , 234–5  
Holmes, Justice O. W. 22  
Hume, D. 6 , 281n.  
Hunter, D. 151n. , 164n.  
Hurford, J. 168  
ignorance 45–50 , 251  
implicature(s) 123 , 142 , 151n. , 164  
inferentialism 122–6  
ethical 124–5  
normative 124–6 *see also* epistemic inferentialism  
informing/informants 44–5 , 60–3 , 74–9 , 82–8 , 127 , 178–91  
indicators 60 , 183–90  
individualism 57–8 , 216

347

James, W. 281n.  
Jaswal, V. K. 152  
justification 1–13 , 51 , 54–7 , 194–201 , 208 , 212–14 , 255 , 278n. , 320 , 323  
Default and Challenge model of 54–67  
externalism about 55 , 63 , 99  
genealogical 140–2  
internalism about 55–7 , 63 , 99n. , 251  
interpersonal 13

personal 12–13  
testimonial 216–40  
Johnson, R. N. 94n.  
Johnston, M. 109n.  
Kallestrup, J. 249n.  
Kant, I. 128  
Kaplan, J. 5 , 80 , 81n.  
Kelly, T. 7 , 11 , 125n. , 282n. , 299–304 , 311n. , 315n.  
knowledge  
concept of 60–2 , 69–76 , 80–8 , 90 , 127  
detached standing 193n.  
externalism about 62 , 263 , 270–2  
genealogical approach to 59 , 63–7 , 81–5 , 140–6  
internalism about 62  
Lewisian theory of 91 , 133–4  
and objectivization 85–8  
practical explication of 73–85  
scientific model of 37  
social theory of 91  
sociology of 52  
value of 89–95 , 100–1 , 104n. , 106–10 , 116  
and well groundedness 64 , 188 *see also* perceptual knowledge  
Koenig, M. 152  
Kolbel, M. 4  
Kornblinth, H. 70n. , 125n. , 148 , 149n. , 194n. , 255n. , 259–60  
K-predicate 79–80  
Krabbe, E. 13n.  
Kripke, S. 70  
Kroll, G. 38n.  
K-state 80  
Kusch, M. 3n. , 70 , 81

Kvanvig, J. 99n. , 101n. , 116n. , 118n.  
Lackey, J. 7 , 191n. , 216n. , 222 , 225n. , 230n. , 231n. , 233–5 , 239n. , 241 , 255n.  
Lane, M. 70n.  
Latour, B. 3–4 , 37 , 39 , 46  
Lasersohn, P. 4–5  
Laudan, L. 27n.  
law 16–22 , 27 , 32–3  
layman vs. expert 52–3  
Lear, L. J. 39 , 48n. , 49n.  
Lehrer, K. 7 , 12 , 257n.  
Leite, A. 240n.  
Levinson, S. 164–6  
Lewis, D. 91 , 109n. , 113n. , 133–40  
Lewisian accounts of social norms 137–40  
Lipton, P. 153 , 255n.  
List, C. 15 , 17 , 25n. , 27  
Locke, J. 114  
logic 13n. , 120–1 , 253  
Lorenzen, P. 13n.  
Ludlow, P. 253n.  
Lynch, M. P. 118n. , 264n.  
Lytle, M. H. 38n. , 48n.  
McDowell, J. 109n. , 160n. , 169n. , 182n. , 184n.  
MacFarlane, J. 4  
McGrath, M. 10  
Mackie, J. L. 115n.  
McNaughton, D. 109n.  
Maher, P. 8n.  
Malmgren, A-S. 194 , 208–11  
memory 29–30 , 200 , 240 , 243–60 , 290–1 , 309–10  
Messing, K. 38–41 , 46–9



meta-philosophy 92  
Mill, J. S. 22  
Millar, A. 116n. , 178n. , 179n. , 180n. , 183n. , 184n. , 190n. , 258n.  
Millikan, R. 154n. , 155–7 , 160n. , 161–3 , 166–7 , 171  
Mills, C. 49n.  
Misak, C. 24  
Moffett, M. 300n.  
moral realism / anti-realism 269  
Moore, C. 152  
Moore, G. E. 287–93 , 297  
Moran, R. 30–4 , 37–9 , 48 , 146 , 177n. , 217n. , 218–26 , 229–40  
Mulhall, S. 182n.  
Nagel, T. 270n.  
natural selection 172–3  
naturalism  
ecological 37  
Quinean 37  
*Naturwissenschaften* 36  
Neander, K. 154n. , 155n.  
Neta, R. 70n. , 113n. , 126n.  
Nichols, S. 70n.  
Nietzsche, F. 66  
normative claims/concepts 2 , 37 , 82 , 114–24 , 127 , 179–80 , 236  
348  
normative commitments 180  
norm(s)  
of argumentation 13n.  
of assertion 90n. , 234–5  
of assurance 235  
of challenge 57–8  
of credulity 221n.

of enquiry 52 , 80  
epistemic 117 , 125 , 129 , 197 , 224n.  
of gender 36  
internalization of 139–46  
moral 129  
of reason 197–8  
of scientific enquiry 40  
social 129–47  
of trust 132–4 , 140–7  
of trustworthiness 132–4 , 140–7  
of truth-telling 130–2 , 136  
of veracity 221n.  
Nozick, R. 17  
Olsson, E. 88n. , 99–103  
Origgi, G. 152n. , 160n. , 166–7  
Owens, D. 216n. , 223n. , 224n. , 233n. , 239n.  
paradox of social epistemic rationality 16–17  
Pargetter, R. 155n.  
Pederson, N. 320n.  
perceptual beliefs 148–9  
perceptual knowledge 149n. , 183–90  
perceptual-recognitional abilities 183–7  
Pettit, P. 15 , 17 , 27 , 135n. , 136–9 , 145  
Piller, C. 106n.  
Pinker, S. 168  
Plantinga, A. 150n. , 195n.  
Plato 89  
politics 30 , 35–8 , 41–9  
Pollock, J. 113n. , 126n.  
Popper, K. 7  
positive epistemic standing 148–9 , 157 , 278n.

postcolonialism 35–8 , 44  
postmodernism 3  
power 52–3 , 234 , 305 , 319  
practical reasoning 15 , 90–1 , 274–7  
practical explication 70–3 *see also* knowledge (practical explication of)  
practices, as essentially rule-governed activities 155 , 179–82 , 259–60  
pragmatism/pragmatists 10 , 86–7 , 116 , 126–28  
presentation as true 203–10 , 249–50  
Price, H. H. 188n.  
Pritchard, D. H. 103–5 , 116n. , 150n. , 178n. , 183n. , 263n. , 273  
Proctor, R. 48n.  
promising 223–6  
prisoner's dilemma 135–7  
Quine, W. V. O. 35–37  
rationality 3 , 4 , 15–17 , 26–7 , 204–10 , 269 , 320n.  
Rawls, J. 24  
reasons 278–96  
epistemic vs. practical 222–42 , 274–7  
recognitional ability 183–92  
Reed, B. 262–3 , 267  
regress objection 238  
Reid, T. 6  
reliabilism/reliabilist theory 31 , 54 , 57 , 99–109 , 127 , 156–7 , 169–74 , 282n.  
relativism 4 , 5 , 129 , 272  
New Age 4 , 5  
Revisionist 4 , 5  
responsible knower 31 *see also* epistemic responsibility  
Ridge, M. 116n. , 119n.  
Rorty, R. 3 , 116 , 127  
Rosen, G. 282n. , 299 , 305n.  
Rosenberg, J. 113n. , 116 , 121n. , 128

Ross, A. 177n. , 217n. , 234–5  
Rysiew, P. 84n. , 113n.  
Sackett, D. L. 43n.  
Sager, L. G. 16  
Sanger, L. 19n.  
Scanlon, T. M. 109n.  
scepticism 35–7 , 55–7 , 67 , 262–3 , 272  
Cartesian 56–9 , 65  
Agrippan 56–8 , 65  
Schaffer, S. 4  
Schiffer, S. 114n.  
Schiebinger, L. 48n.  
Schmitt, F. 15 , 216n.  
Schroeder, M. 114n. , 120n.  
Schwartz, P. 157  
Sellars, W. 55 , 113n. , 120n. , 121n. , 125  
Seller, A. 46n.  
sensibility 182–8  
semantics/semantical claims, theses or concepts 5 , 114–16 , 119n. , 123n. , 127–8  
Shaked, M. 8n.  
Shapin, S. 3–4  
Shope, R. 94n. , 109n.  
silence 35–7 , 243–61  
sincerity 140–5 , 189–90 , 218–20  
Skyrms, B. 8  
Smith, M. 94n. , 109n. , 115n.  
Smith, M. B. 38n. , 40n.  
Sober, E. 162n.

social epistemology 1 , 2 , 30 , 36 , 51 , 91 , 113 , 116 , 126–8 , 256  
expansionist 15–25  
naturalized 44  
preservationist 5–15  
revisionist 3–5  
social identity 52–3  
social imaginery 33–5 , 44–9  
socialization 260  
social systems 18–26 , 72–3 , 169  
Sosa, E. 7 , 31n. , 113n. , 150n. , 195n. , 263n. , 312  
Sperber, D. 152n. , 160n. , 161 , 162n. , 166–7 , 171–2  
Sripada, C. S. 139n.  
Stanley, J. 10 , 113n. , 126n.  
State of Nature see epistemic State of Nature  
Steup, M. 31n.  
Stevenson, C. 115n.  
Stich, S. 70n. , 139n.  
Stroud, B. 263n.  
subjugated knowledges 48  
supervenience 93  
Sullivan, S. 48n.  
Swamping Problem 89–111  
Talisse, R. 24  
Tanesini, A. 63n.  
time 53 , 54 , 65–8 , 290–1  
telling, as a speech act 131 , 175–92 , 218n. , 221n. , 223n.  
testimony 6 , 19–20 , 29–33 , 38–49 , 140–7 , 175–92 , 149–53 , 171 , 194 , 199–203 ,  
208 , 212 , 243–55 , 260–1  
assertion view of 239–41  
assurance view of 216–42  
agreement version 217n. , 220–1

communicative norm version 220–1  
beliefs based on 149 , 151n.  
belief expression model of 239n.  
defaultist view of 188n. , 191  
statement view of 239n.  
straightforward cases of 175–7 , 188–93  
Thompson, D. 24  
Timmons, M. 116n.  
Tomasello, M. 166n , 171  
Toulmin, S. 13n.  
Travis, C. 190n.  
Trout, J. 70n.  
trust 78–9 , 217n. , 226–41 see also norm(s) (of trust)  
trustworthiness 11 , 33 , 188–90 , 218n. , 219–20 see also norm(s) (of trustworthiness)  
truth 76–8 , 264n.  
truth-telling 130–2 , 136 , 162n.  
Tuana, N. 48n.  
Tuomela, R. 15  
Unger, P. 12n. , 113n. , 148 , 223n.  
uniformity 302–5 , 321  
value 23 , 140–5  
extrinsic 103–6  
final 103–11  
instrumental 116 , 141–2  
intrinsic 103–6 , 141–5  
Moorean conception of 104–5  
as response-dependent 107–10 see also knowledge (value of)  
Van Inwagen, P. 282n. , 299n. , 302n. , 304  
Wagner, C. 7  
Warfield, T. 7  
warrant 148–58 , 169 , 173 , 194–203 , 208–13 , 271

a priori 198–9 , 205–12

preservation of 199–201 , 212

Wedgwood, R. 123n. , 299

Weinburg, J. 70n.

Welbourne, M. 31–3 , 37 , 48 , 60n. , 179n. , 188n.

White, R. 279n.

Wiggins, D. 109n.

Williams, B. 60n. , 66n. , 70 , 81 , 131 , 140–7

Williams, M. 12 , 54–9 , 61–8 , 113n.

Williams, P. J. 29 , 32–4 , 37–8 , 46 , 49

Williamson, T. 12n. , 90n. , 113n. , 126n. , 179n. , 223n.

Wittgenstein, L. 12 , 32 , 47 , 70

Woods, J. 13n.

Woolgar, S. 3

Wright, C. 4–5 , 271

Wright, L. 154n.

Zagzebski, L. 113n.