

OXFORD

THE EPISTEMIC LIFE OF GROUPS

Essays in the Epistemology of Collectives

edited by MICHAEL S. BRADY & MIRANDA FRICKER



The Epistemic Life of Groups

MIND ASSOCIATION OCCASIONAL SERIES

This series consists of carefully selected volumes of significant original papers on predefined themes, normally growing out of a conference supported by a Mind Association Major Conference Grant. The Association nominates an editor or editors for each collection, and may cooperate with other bodies in promoting conferences or other scholarly activities in connection with the preparation of particular volumes.

Director, Mind Association: Julian Dodd

Publications Officer: Sarah Sawyer

Recently Published in the Series:

Reality Making

Edited by Mark Jago

The Metaphysics of Relations

Edited by Anna Marmodoro and David Yates

Thomas Reid on Mind, Knowledge, and Value

Rebecca Copenhaver and Todd Buras

The Highest Good in Aristotle and Kant

Joachim Aufderheide and Ralf M. Bader

Foundations of Logical Consequence

Edited by Colin R. Caret and Ole T. Hjortland

The Highest Good in Aristotle and Kant

Edited by Joachim Aufderheide and Ralf M. Bader

How We Fight: Ethics in War

Edited by Helen Frowe and Gerald Lang

The Morality of Defensive War

Edited by Cécile Fabre and Seth Lazar

Metaphysics and Science

Edited by Stephen Mumford and Matthew Tugby

Thick Concepts

Edited by Simon Kirchin

The Epistemic Life of Groups

*Essays in the Epistemology
of Collectives*

Edited by
Michael S. Brady and Miranda Fricker

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

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

© the several contributors 2016

The moral rights of the authors have been asserted

First Edition published in 2016

Impression: 1

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, by licence 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 work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2015946838

ISBN 978-0-19-875964-5

Printed in Great Britain by
Clays Ltd, St Ives plc

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

<i>List of Contributors</i>	vii
Introduction <i>Michael S. Brady and Miranda Fricker</i>	1
Part I. Epistemology	
1. Mutuality and Assertion <i>Sanford C. Goldberg</i>	11
2. Fault and No-Fault Responsibility for Implicit Prejudice: A Space for Epistemic ‘Agent-Regret’ <i>Miranda Fricker</i>	33
3. On Knowing What We’re Doing Together: Groundless Group Self-Knowledge and Plural Self-Blindness <i>Hans Bernhard Schmid</i>	51
Part II. Ethics	
4. The Social Epistemology of Morality: Learning from the Forgotten History of the Abolition of Slavery <i>Elizabeth Anderson</i>	75
5. Group Emotion and Group Understanding <i>Michael S. Brady</i>	95
6. Changing Our Mind <i>Glen Pettigrove</i>	111
Part III. Political Philosophy	
7. The Epistemic Circumstances of Democracy <i>Fabienne Peter</i>	133
8. The Transfer of Duties: From Individuals to States and Back Again <i>Stephanie Collins and Holly Lawford-Smith</i>	150

9. Four Types of Moral Wriggle Room: Uncovering Mechanisms of Racial Discrimination <i>Kai Spiekermann</i>	173
Part IV. Philosophy of Science	
10. Collective Belief, Kuhn, and the String Theory Community <i>James Owen Weatherall and Margaret Gilbert</i>	191
11. Collaborative Research, Scientific Communities, and the Social Diffusion of Trustworthiness <i>Torsten Wilholt</i>	218
<i>Bibliography</i>	235
<i>Index</i>	251

List of Contributors

Elizabeth Anderson	University of Michigan, Ann Arbor
Michael S. Brady	University of Glasgow
Stephanie Collins	University of Manchester
Miranda Fricker	University of Sheffield
Margaret Gilbert	University of California, Irvine
Sanford C. Goldberg	Northwestern University
Holly Lawford-Smith	University of Sheffield
Fabienne Peter	University of Warwick
Glen Pettigrove	University of Auckland
Hans Bernhard Schmid	University of Vienna
Kai Spiekermann	London School of Economics
James Owen Weatherall	University of California, Irvine
Torsten Wilholt	Leibniz Universität Hannover

Introduction

Michael S. Brady and Miranda Fricker

Groups engage in epistemic activity all the time—whether it be the active collective inquiry of localized epistemic communities such as scientific research groups or crime-detection units, or the heavily institutionally structured evidential deliberations of tribunals and juries, or the more spontaneous and imperfect information-processing of the voting population. In the philosophy of mind and action there is a matured literature advancing competing theories of what groups are and how they do what they do.¹ Such debates principally turn on whether groups are best construed as no more than the sum of the individuals that compose them acting in concert, or whether certain forms of practical and intentional interdependencies suffice to generate a new agent, a distinct group agent that is greater than, or at least different from, the sum of its parts. More recently, social epistemology has also flourished, expanding and making connections with other areas of philosophy.² With these two philosophical lenses now beginning to align themselves, the new vista before us is that of *collective epistemology*—a natural next step for social epistemology.³

One of the key aspects of group agency is the deliberation that lies behind it, and the various epistemic commitments and capacities that are involved in such deliberations. The relevant debates in epistemology, concerning such things as competing accounts of collective belief, justification, and knowledge, have now begun to flourish. The

¹ A few indicative landmarks in this expansive territory might be Gilbert (1989 and 2000); Bratman (1999); List and Pettit (2011); Tuomela (2013).

² See e.g. recent edited collections in social epistemology: Goldman and Whitcomb (2010); Haddock, Millar, and Pritchard (2010). Also Grasswick and Webb (2002), which was part of an issue of *Social Epistemology* devoted to the theme of 'Feminist Epistemology as Social Epistemology'; in which general connection, see also Grasswick (2013).

³ To track the rise of 'collective epistemology', see Gilbert (2004), and Tollefsen (2007a); the special issue of *Social Epistemology* edited by Mathieson (2007); and most recently collections edited by Schmid, Sirtes, and Webe (2011) and Lackey (2014).

essays in this book, however, are not on the whole directly focused on these issues, but rather explore different epistemic aspects of the behaviour of different sorts of group—institutional bodies, the moral community, informal groups, religious communities, the state, or simply the population at large. To this extent the book is part of an expansionist trend in epistemology of the last decade, consisting in the exploration of new epistemological projects that go beyond the traditional problems such as the refutation of scepticism, the nature of warrant, or the analysis of knowledge. The essays collected here explore different aspects of the epistemic practices of groups, thereby indicating the great range of ways in which epistemological issues permeate the well-functioning, or otherwise, of different kinds of human collectivity.

This volume collects essays by leading philosophers of different but coinciding persuasions in order to generate a more multilateral conversation than is ordinarily possible—but which is highly desirable given the manifestly common concerns, and the breadth of significance associated with these matters of collective epistemic practice. The essays each explore some region of our collective epistemic practice; and each essay has a particular focus that brings it under one of the following broad section headings: Epistemology, Ethics, Political Philosophy, and Philosophy of Science. The essays in the Epistemology section (Part I) address topics that make them fit squarely within the core social epistemological remit; the essays in the other sections address epistemological strands running through topics that primarily belong in other subject areas of philosophy. Together the essays indicate the richness and complexity of the philosophical issues generated by thinking about the epistemic life of groups. In what follows we give a brief outline of the individual chapters, before closing with remarks about some of the central issues raised.

Part I. Epistemology

Sanford C. Goldberg's chapter, 'Mutuality and Assertion', illustrates the importance of collective epistemology to a core topic in recent epistemology, namely epistemic disagreement. Goldberg's central question is this: can assertions be warranted under conditions of systematic disagreement—disagreement of the sort we find in philosophy, politics, religion, and the more theoretical parts of the social and natural sciences? On the one hand, he argues, there are strong reasons to regard assertion as governed by a demanding epistemic norm (such as knowledge), and it is plausible to think that in contexts of systematic peer disagreement we sometimes (often?) fail to attain knowledge. On the other, the practice of assertion persists in these areas, even under conditions of systematic peer disagreement. Indeed, the discipline of philosophy would appear to depend on this practice. Goldberg proposes that this tension can be resolved by appeal to

the hypothesis that the standard set by assertion's norm is fixed in part in terms of something beyond the individual: namely, what is mutually believed by the speaker and her audience in the context in which the assertion is made. This is what he terms the *Mutual Belief Norm*. His chapter aims to provide independent grounds for this hypothesis. His central argument is that we can do so by appeal to Grice's guiding idea that conversation is a cooperative activity between agents, so that the rationality of speech contributions is to be understood by reference to the group context in which they are operating.

Miranda Fricker's chapter, 'Fault and No-Fault Responsibility for Implicit Prejudice: A Space for Epistemic "Agent-Regret"', explores different forms of epistemic responsibility, individual and collective, as regards the influence of prejudice on judgement. On the whole, if one makes judgements that are significantly influenced by prejudicial bias, then one is epistemically at fault, so that epistemic blame would be justified, including self-blame. What about cases where the prejudice in question is an 'implicit bias' (non-conscious, automatic, evidence-resistant, and possibly contrary to one's beliefs)? Here too, Fricker argues, the default is that we stand as blameworthy, though allowing that there may often be extenuating circumstances that diminish the degree of appropriate blame. Compare an entrenched character trait of which the subject is unaware partly because it does not fit with her self-conception—'implicit' selfishness, perhaps.

Fricker asks, however, whether there are circumstances in which we are guilty of implicit prejudice and yet where blame *is* entirely inappropriate (not merely extenuated). An example might be a case of *environmental epistemic bad luck*: where there is prejudice in the epistemic environment, *and* one has no reason to suspect that this is so, resulting in an *epistemically innocent* inheritance of environmental prejudice. Where this is so, argues Fricker, we confront the space of a 'no-fault epistemic responsibility'—the epistemic analogue of 'agent-regret'. We are not epistemically blameworthy, and yet we still have responsibility, as is revealed in the fact that there are epistemic obligations which apply to us specifically because it was through *our* epistemic agency that the prejudiced judgement was made. The fulfilment of those epistemic obligations, it is argued, will typically require the promotion of collective institutional measures to militate against the influence of prejudice in the future. This will be the case insofar as the practical remit of the individual's obligation coincides with an existing responsibility of the institution—for instance, where an individual employee has responsibility for a given promotions process with the organization. This coincidence of areas of responsibility means that the individual's responsibility to ameliorate a situation of potential bias quickly generates a collective, organizational one. And so, Fricker concludes, any counter-biasing epistemic responsibilities of

individuals acting under the auspices of an institutional body (as an employee, for instance) will tend to beget collective epistemic responsibilities to take ameliorative counter-biasing measures.

Hans Bernhard Schmid contributes the third chapter to this section: ‘On Knowing What We’re Doing Together: Groundless Group Self-Knowledge and Plural Self-Blindness’. Schmid is concerned with whether an influential view about intentional behaviour at the individual level is also true at the collective level. This is the view that in order to act intentionally, an agent needs to know what she is doing, where this knowledge is ‘groundless’—that is, non-observational and non-inferential. His central question is this: is our knowledge of what we are doing *together with others* of the same groundless kind? Schmid begins by highlighting the central features of groundless self-knowledge on Anscombean lines: these are first-person identity, first-person perspective, first-person commitment, and first-person authority. He then considers an argument that such knowledge is not available at the group level, on the grounds that a requirement for group knowledge is that each individual needs to know what her partners are doing, and this would seem to require observation and inference. Schmid concludes, however, that a moderate version of the claim that there is groundless group self-knowledge is defensible, if we maintain that the relevant knowledge involves plural pre-reflective and non-thematic self-awareness of what it is that the group members are jointly doing. Thus there is indeed a common structure to individuals’ and groups’ knowledge of their actions.

Part II. Ethics

Elizabeth Anderson’s chapter, ‘The Social Epistemology of Morality: Learning from the Forgotten History of the Abolition of Slavery’, opens the section on Ethics. Anderson’s concern is how social groups learn from history, and how the organization of social groups bears on the prospects for improvements to group beliefs. Anderson’s focus is on the particular history of slavery. She notes that during the nineteenth century, the belief that individuals have a right against being enslaved became a nearly worldwide consensus. Most people today believe that this change in moral convictions was a case of moral learning. But Anderson is concerned with how we can know this, or similar claims about moral progress, without begging the question in favour of our current beliefs. She proposes to answer this question by developing a naturalized, pragmatist moral epistemology through case studies of moral lessons people have drawn from the history of abolition and emancipation. Anderson argues that processes of contention, in which participants challenge existing moral and legal principles governing

interpersonal relationships, play critical roles in moral learning. Contention may take the form of argument, but it takes many other forms as well, including litigation, protest, and revolution. Anderson concludes that progress in moral inquiry requires that groups are receptive to and open to the perspectives of others, and not simply of those in authority; it requires ‘the practice of epistemic justice by and for all’.

Collective understanding is the focus of Michael Brady’s contribution to this volume, ‘Group Emotion and Group Understanding’. Brady aims to explain how group emotion can have positive epistemic value in so far as it promotes group understanding; and he argues that this epistemic good would be very difficult to achieve, in many cases, without group emotion. After outlining philosophical, psychological, and neuroscientific support for the view that individual emotion promotes individual understanding, Brady applies this picture to group emotion and group understanding, and illustrates the connection between the two by focusing on the phenomenon of *public inquiries*. On this view, these inquiries are both motivated by group emotion, and aim at the attaining and dissemination of information so that some group understands what has happened. Without group emotion, he argues, it is unlikely that this collective epistemic good would be achieved. If this is correct, then group emotion is more valuable, from an epistemic standpoint, than traditional thinking on this issue supposes.

In ‘Changing Our Mind’, Glen Pettigrove considers the ways in which groups revise their beliefs, and proposes an account of group belief revision that allows modifications along a number of different dimensions. In particular, Pettigrove proposes an account of group belief revision that can accommodate modifications of (i) propositional content, (ii) non-propositional content, (iii) understanding, and (iv) conception. He develops his account by focusing on communities that are less discussed in the literature on social epistemology, namely moral and religious communities. By focusing on these communities, Pettigrove argues that Margaret Gilbert’s account of group belief revision needs to be supplemented: while the view is adequate to changes in collective propositional knowledge-that, it fails to capture or accommodate revision of collective holistic knowledge-that, and in particular cannot accommodate revision to belief in groups such as religious communities or charitable organizations, which are built around normative commitments.

Part III. Political Philosophy

Fabienne Peter’s chapter, ‘The Epistemic Circumstances of Democracy’, focuses on a long-standing question in political philosophy, namely: ‘does political

decision-making require experts, or can a democracy be trusted to make correct decisions?’ Peter notes that the traditional debate about this issue is *instrumentalist*, in so far as it is thought that the determining factor for the legitimacy of political institutions is the epistemic status of the outcomes of decision-making processes. Supporters of democracy argue that this system produces optimal outcomes and hence can be trusted, whilst critics of democracy argue that outcomes of democratic decision-making are sub-optimal and hence the system cannot be trusted. Peter argues—against the instrumentalist approach—that attempts to defend democracy on epistemic instrumentalist lines are self-undermining. She proceeds to develop an alternative, procedural, epistemic defence of democracy, arguing that there is a *prima facie* epistemic case for democracy whenever there is no procedure-independent epistemic authority available on the issue that is to be decided.

In ‘The Transfer of Duties: From Individuals to States and Back Again’, Stephanie Collins and Holly Lawford-Smith consider whether a standard model for the transfer of duties from individuals to collectives applies to states’ duties. On the standard model, individuals sometimes pass their duties on to collectives, which is one way in which collectives can come to have duties. This involves certain crucial epistemic transactions: notably, that the individual communicate to the collective the knowledge that she wills the transfer of duty; and that the individual makes all reasonable efforts to bring it about that she has a reasonable belief that the collective will indeed discharge the duty appropriately. The collective discharges its duties by acting through its members, which involves distributing duties back out to individuals. Individuals put duties in and get (transformed) duties out. But can this general account make sense of *states’* duties? And if so, to what extent? Do some of the duties we typically take states to have come from individuals having passed on certain individual duties? The authors note that there are complications to the picture: states can discharge their duties by contracting fulfilment out to non-members; states seem able to dissolve the duties of non-members; and some duties of states are not derived in this way. They argue that these complications do not undermine the general account and its application to states. Furthermore, Collins and Lawford-Smith show that the application has an interesting upshot: by asking which individuals *robustly* participate in this process of duty transfer-and-transformation with a given state, they show how we can begin to get a grip on who counts as a *member* of that state.

A different political problem is taken up by Kai Spiekermann in his chapter, ‘Four Types of Moral Wriggle Room: Uncovering Mechanisms of Racial Discrimination’. Spiekermann describes recent experiments in behavioural economics

which reveal that individuals frequently use so-called ‘moral wriggle room’ to avoid complying with costly normative demands. ‘Wriggle room’ describes our tendency to shape and modify our belief-system so as to convince ourselves that our behaviour is morally appropriate or satisfactory, and to thereby avoid stringent moral obligations. Spiekermann develops a typology of ‘moral wriggle rooms’ that helps to illustrate different opportunities for strategic information manipulation, and shows how moral wriggling can often operate in an unconscious, yet systematic way. He then notes that the experimental literature tends to be focused on individual behaviour; however, failures to meet obligations of inquiry are rooted in social practices and institutions. For example, one’s individual ability to maintain biased beliefs is much higher when all of one’s peers have the same biased beliefs. As a result, problems generated by moral wriggling must be addressed at the level of the group. Spiekermann illustrates this issue for social moral epistemology using the case study of racial discrimination, and considers a number of options we might employ to mitigate this problem.

Part IV. Philosophy of Science

James Owen Weatherall and Margaret Gilbert’s chapter, ‘Collective Belief, Kuhn, and the String-Theory Community’, opens the final section of the volume, in which scientific aspects of collective epistemology are discussed. The chapter begins with Gilbert’s well-known account of collective belief: this involves a joint commitment of certain parties, who constitute what she refers to as a plural subject. Gilbert has, in previous work, argued that ascriptions of beliefs to scientific communities commonly involve appeal to collective belief understood in her sense. This raises a potential problem when it comes to scientific change, however. For if Gilbert’s view of collective belief as involving joint commitment is correct, and some of the belief ascriptions in question are true, then the members of some scientific communities have obligations that may act as barriers both to the fair evaluation of new ideas and to changes in scientific consensus. The authors argue that this may help to explain Thomas Kuhn’s observations on ‘normal science’, and go on to develop the relationship between Gilbert’s proposal and several features of a group of physicists working on a fundamental physical theory called ‘string theory’, as described by physicist Lee Smolin. Weatherall and Gilbert argue that the features of the string theory community that Smolin cites are well explained by the hypothesis that the community is a plural subject of belief. As a result, reflection on the practices of an actual scientific community provides further support for Gilbert’s account of group belief.

The final chapter focuses on issues of trustworthiness in scientific research. In 'Collaborative Research, Scientific Communities, and the Social Diffusion of Trustworthiness', Torsten Wilholt argues for the thesis that when we trust the results of scientific research, that trust is inevitably directed, at least in part, at collective bodies rather than at single researchers. As a result, he proposes that reasonable assessments of epistemic trustworthiness in science must attend to these collective bodies. Wilholt supports his thesis by first explaining the collaborative nature of most of today's scientific research. He argues that the trustworthiness of a collaborative research group does not supervene on the trustworthiness of its individual members, and points out some specific problems for the assessment of epistemic trustworthiness that arise from the specific nature of today's collaborative research. Wilholt then argues that the social diffusion of trustworthiness goes even further; on his view, we always also need an assessment of the trustworthiness of the respective research community as a whole. Communities, he claims, play an essential role in the epistemic quality management of science. To see why this role is indispensable, Wilholt investigates and delineates three dimensions of what is desirable in a method of inquiry: the reliability of positive results, the reliability of negative results, and the method's power. Every methodological choice involves a trade-off between these three dimensions. The right balance between them depends on value judgements about the costs of false results and the benefits of correct ones. Conventional methodological standards of research communities impose constraints on these and thereby harmonize the implicit value judgements. Trusting that the research community has done this in a suitable way is thus always part of placing our trust in a scientific result.

Together, these studies indicate the extraordinary reach and internal diversity of the emerging field of Collective Epistemology. The strange individualism of English language philosophy is a historical contingency with a complex aetiology, and it has impacted on many different areas of philosophical discussion in ways that are not always easy to diagnose. But, in addition, it has also simply concealed many important philosophical questions from view, and one such is the range of questions addressed by the authors contributing to this book—questions relating to collective epistemic practice and its significance for how we know everyday things, how we pursue a shared moral life, how we conduct ourselves in professional contexts, how we design and modify our political institutions, and the forms of inquiry that govern the advancement of science. There are, of course, many other possibilities for collective epistemology than are exemplified here. Our hope is that this book will help further extend philosophy's conception of the proper domain of epistemology, thereby opening up many new promising avenues of understanding.

PART I

Epistemology

1

Mutuality and Assertion

Sanford C. Goldberg

1. The Problem: Assertoric Practice and Assessment

The aim of this chapter is twofold. First, I aim to develop a context-sensitive account of the norm of assertion in which the requirement set by this norm, taken to be epistemic in its content, crucially involves the mutual expectations of interlocutors within an epistemic community. Second, I aim to provide independent grounds for thinking that this account is true, and I aim to defend it against various objections.

Before doing either of these things, I want to present a problem, to which such an account (if independently motivated) would be a solution. The problem is that there are two background assumptions, each plausible in its own right (if not universally endorsed), yet which are incompatible with one another.

The first assumption concerns the norm of assertion itself. It is plausible to think that assertion has an epistemic norm E: speaker S's assertion that p is proper only if S satisfies E with respect to p, i.e. only if S has the relevant warranting authority regarding p. What is more, it is plausible to think that E is a demanding epistemic standard. The most prevalent view in the literature is that it is knowledge. If this is correct, then the norm of assertion tells us this: one should not assert that p, unless one knows that p.¹ But whether or not one thinks

I would like to thank Miranda Fricker and Michael Brady for very helpful comments on earlier versions of this chapter. I would also like to thank audiences at the various places at which I have given this paper: the philosophy departments at Monash University, the University of New South Wales, LOGOS (University of Barcelona), Vanderbilt University, the University of Cologne, and the University of Aarhus; and also the 2012 Orange Beach Epistemology Workshop on Social Epistemology, the 2011 Kentucky Philosophical Association meeting, and the 2011 Conference on 'Collective Epistemology: The Epistemic Life of Groups', at the Institute of Philosophy, University of London.

¹ This view has been endorsed by Unger (1975), Williamson (1996), DeRose (1996), Hawthorne (2003), and Stanley (2005).

that knowledge is the norm, it is widely agreed that the standard itself, E, is demanding.

The second assumption concerns the practice of assertion. In particular, there are plenty of cases in which, whether owing to pervasive disagreement (philosophy, politics, theoretical areas in the social and natural sciences, etc.) or low epistemic expectations (difficulties of achieving *knowledge* in highly theoretical areas), few if any speakers have any epistemic credentials such as knowledge, or knowledge-sufficient justification or warrant, to make assertions in these areas. And yet assertions continue to be made, unabated, in these areas. Philosophy is a particularly interesting domain in which to develop this problem. Elsewhere² I have argued that disagreement in philosophy is a special case of peer disagreement, where, whatever one thinks about the epistemic significance of disagreement as such, the sort of *systematic* disagreements one finds in philosophy make the prospects for epistemically high-grade belief remote.³ Insofar as there is no epistemically high-grade belief at all under these conditions—a conclusion for which I argue in Goldberg (2013a)—I will call these conditions the conditions of *diminished epistemic hopes*. Yet assertions continue in philosophy, even under conditions of systematic disagreement, and hence diminished epistemic hopes. Indeed, the practice of philosophy would appear to depend on this.⁴

Taken together, these two assumptions suggest an unhappy conclusion: insofar as the practice of assertion continues even under conditions of diminished epistemic hopes, these assertions are *systematically improper*. In short, we appear to face a stark choice: accept the unhappy conclusion, or else reject one of the two assumptions. Yet none of these options seems particularly happy.

Consider first the option to accept the unhappy conclusion. We then face a situation in which every assertion made under conditions of diminished epistemic hopes is improper. If I am right that the practice of philosophy depends on the continued making of assertions under such conditions, then we would reach the unhappy conclusion that the practice of philosophy is at least to this extent suspect. Perhaps we might grant that, though assertions in question are improper, we let each other ‘get away with’ them anyway. But this suggestion is unhelpful: why would we think any higher of philosophy merely because we let each other get away with shoddy assertions? On the contrary, this would seem to condemn the practice twice-over: first, for improper assertions, second, for a

² Goldberg (2013a).

³ By ‘epistemically high-grade belief’ I mean belief that amounts to knowledge or else is based on a knowledge-sufficient justification.

⁴ See Goldberg (2013a and 2013b).

refusal to acknowledge this impropriety.⁵ And what goes for philosophy goes for other areas in which the practice of assertion continues even under conditions of diminished epistemic hopes.

Consider next the option to reject the first assumption, to the effect that assertion has a demanding epistemic norm. To be sure, there are people who argue that assertion has no norm at all,⁶ and others who argue that while assertion has a norm the norm is not epistemic.⁷ Still, the vast majority of people who work on assertion seem to regard it as having an epistemic norm of some sort or other. Insofar as they are right, rejecting the hypothesis that assertion has an epistemic norm is already a cost. But now consider what must be done if, having accepted that assertion has an epistemic norm of some sort or other, we still want to reject the first assumption above. The conclusion would have to be that the norm is not particularly demanding. We would then have to replace the assumption of a *demanding* epistemic norm with an alternative account—one whereby the epistemic standard, E, is no more demanding than is needed in order to represent assertoric practice in these areas as desired. To be sure, there *are* unwarranted assertions under conditions of diminished epistemic hopes; the point above is rather that *not all* assertions made under these conditions are improper (i.e. some are proper). But even to get this result we would need to weaken E quite a bit. For if I am right that (in contexts of systematic peer disagreement) we face a situation of diminished epistemic hopes, then E would have to be weakened to the point where it demands something weaker than justified (or rational) belief. Of course, no sooner do we do this, than we have an obviously overly-permissive norm—one that sanctions a good deal of assertions we would want to regard as improper.

It is perhaps worthwhile discussing the sort of overly-permissive norm that emerges from our dialectic so far, if only to see why norms in this class would appear to be non-starters. To this end, consider the sort of norm that might be satisfied in conditions of diminished epistemic hopes. One candidate norm is that of mere belief. If the norm of assertion is mere belief, then a speaker might well satisfy it under conditions of epistemically diminished hope. But there would appear to be something curious in the idea that an assertion is warranted (proper, permissible) so long as the speaker *believes* what she says. For in effect this collapses (something like) the sincerity condition on assertion with assertion's

⁵ Although the assertoric practice of philosophy is illustrative, it is not unique in this respect. Consider the theoretical parts of the social and natural sciences.

⁶ See Pagin (2011). See also MacFarlane (2011) for an inventory of the sorts of views one can take about the nature of assertion.

⁷ See Weiner (2005).

norm, and such a collapse would appear to have unacceptable implications. For example, it would follow that the incompetent believer is nevertheless a proper asserter (so long as he asserts only what he believes). One might think to rectify this by saying that it is a deficiency in assertion if one expresses what in fact is a deficient (unwarranted, unjustified) belief; such a view is presented by Bach (2008). But we would want to know why this is so: if an assertion is proper so long as the speaker believes the asserted content, it is unclear how the epistemic deficiency of the belief itself should affect the propriety of the assertion. After all, not all of the ways in which we would criticize an assertion bear on its propriety *qua* assertion: assertions that are rude, irrelevant, or made in too loud a voice are thereby criticizable, but if the assertion in question expressed the speaker's knowledge, then, at least with respect to the standards of the kind of speech act it was, the assertion was proper. Thus the proponent of the belief norm would appear precluded from saying that the assertion of an unjustified belief is improper; such a proponent is restricted to saying that the assertion is criticizable on *other* grounds, in violation of a standard other than that pertaining to the kind of speech act it was. This appears to me to be a great cost. Of course, what goes for the norm of belief also goes, *mutatis mutandis*, for other candidate norms that can be satisfied under conditions of epistemically diminished hopes. Such norms are too permissive.

But if we reject the option to weaken the norm of assertion to accommodate the cases in question, and if we also reject the option of accepting the unhappy conclusion that there are no proper assertions under conditions of epistemically diminished hope, we are left with the option to reject the second assumption above. This was the assumption that there are many areas in which the making of assertions continues even under conditions of systematic disagreement, and so even under conditions in which there is a recognition that the hope for knowledge and justified belief is remote.⁸ To reject this assumption is to hold that assertions are not (typically, standardly) made in these circumstances after all. Why might one think this? Well, one might argue that what appear to be assertoric speech acts *aren't really assertions after all*. Perhaps they are hedged or qualified assertions; or perhaps they are a different speech act altogether. The difficulty with this reaction is that it seems seriously revisionary: although there surely *are* hedged or qualified assertions made in these contexts, it is also the case that (to all outward appearances) there are many examples of *straight assertions* in these contexts as well. To reject the second assumption is to have to treat these

⁸ An initial reaction to this assumption might be to accept it, but hold that all such assertions are improper. This, in effect, would be to accept what I called the unhappy conclusion.

as cases in which things are not as they seem. I do not claim that this is unacceptable; only that, since it is revisionary, it would be better to see if we can address the problem without having to go this route.

With this in mind, I propose the following *desiderata* for a solution to our problem. We would like a solution to satisfy the following three conditions:

- (1) it should not weaken E—the epistemic standard provided by the norm of assertion—to the point of irrelevance or insignificance;
- (2) it should recognize the strengths of the arguments made on behalf of thinking that E is *knowledge*, or some other substantial epistemic property (such as knowledge-sufficient justification); and
- (3) it should enable us to regard the relevant class of assertions as broadly proper despite the fact that few if any speakers have any substantial epistemic credentials to make assertions in these areas.

Now, it might be thought that there is no solution that satisfies all of (1)–(3). If this thought were true, then we would need to revisit the other options (dismissed above). However, in what follows I will argue that there is a solution that satisfies all of (1)–(3); and I will argue as well that this solution enjoys support independent of its offering a solution to the present problem.

2. Toward a Solution

The problem we are considering pertains to any domain in which the making of assertions continues despite diminished epistemic hopes. Above I mentioned one such domain—that of philosophy, where it is an ordinary part of the practice to make assertions under these conditions. I noted that it is tempting to describe this part of philosophical practice by saying that participants in conversations regularly let each other ‘get away with’ the assertions in question. Again, if this is merely a matter of letting one another ‘get away with’ what in fact are improper assertions, this is doubly bad for the practice. But we might well wonder whether this is the most perspicuous description of the practice. Perhaps it is not a matter of letting each other ‘get away with’ anything; perhaps the assertions in question are proper after all.

Of course, anyone who would like to try to make out such a view, and so who aims to satisfy *desideratum* (3) above, must face an immediate question: how to square the hypothesis that these assertions are proper after all, with the claim that few if any speakers in these areas enjoy any substantial epistemic standing with respect to the propositions they are asserting? In short, it can seem that the move to satisfy (3) will come at the cost of having to weaken the norm’s standard to the

point of irrelevance, and hence at the cost of failing to satisfy *desideratum* (2) above. I should emphasize as well that this difficulty has nothing in particular to do with assertoric practice in philosophy: it is a problem facing *any* discourse in which assertions continue to be made (as an ordinary part of the discourse) under conditions of diminished epistemic hopes.

In this section I argue that the difficulty can be met. I propose to do so by developing an account of assertion on which the standard provided by assertion's norm is set in a context-sensitive way. In this section I motivate such an account by appeal to broadly Gricean considerations; in the section following I argue that the account delivers the desired outcome, satisfying *desiderata* (1)–(3) above.

An account of assertion is an account of a type of speech act, and as such it should be presented against the background of our best understanding of speech acts. Paul Grice (1968/89) has provided an important grounding principle for such an understanding. Regarding speech as a rational, cooperative activity, Grice formulated the familiar Cooperative Principle as capturing a core part of the rationality of particular acts of this sort:

Cooperative Principle (CP)

Make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

(Grice 1968/89, 26)

Now most people who have employed Grice's CP (or the other elements of Grice's picture) have done so with an eye on characterizing the *content* dimension of communication. That is, they use Grice's framework to provide an account of how speakers manage to communicate more than they (strictly and literally) say, and of how hearers manage to recover what is communicated when this goes beyond what is (strictly and literally) said. But I see no reason why we can't use Grice's insight to shed light on the dimension of (illocutionary) *force*. This is what I propose to do.

One might wonder how Grice's CP can shed any light on the dimension of illocutionary force in general, and on assertion in particular. But this suggestion is not as outlandish as one might suppose. On the contrary, it is a natural one. To see this, consider that Grice went on to present various maxims that he regarded as falling out of CP. One of these was the maxim he labeled 'Quality', which I repeat here:

Quality (Q)

Do not say what you believe to be false. Do not say that for which you lack adequate evidence. (Grice 1968/89, 27)

Now it is true that Grice's maxims are explicitly aimed at characterizing the notion of *saying* something. Still, when it comes to Q itself, it is not a far stretch to regard the maxim as contributing to our understanding of the more specific speech act of assertion. Indeed, we might well think that the 'quality' dimension of Grice's notion of saying just is a proposed characterization of the norm of assertion. On this picture, the speech act of assertion is governed by two rules: you shouldn't assert what you believe to be false, and you shouldn't assert that for which you lack adequate evidence. But precisely when is one's evidence 'adequate'? I submit that we should answer this question by appeal to the CP itself: the standards for adequacy of evidence are determined, at least in part, by 'the accepted purpose or direction of the talk exchange in which you are engaged'. If this is correct, then we have characterized a feature of the illocutionary force of an assertion by appeal to CP.

Still, we need an account of how 'the accepted purpose or direction of the talk exchange in which you are engaged' might serve to fix the standards of evidential adequacy. I propose that we can develop such an account in terms of Bach and Harnisch's useful (1979) notion of *mutual belief*. The following is their gloss on the role that mutual belief plays in the sort of inferences that are made in the course of the production and comprehension of speech:

Mutual Belief

If *p* is mutually believed between S and H, then (1) not only do S and H believe *p*, but (2) each believes that the other takes it into account in his thinking, and (3) each, supposing the other to take *p* into account, supposes the other to take him to take it into account.

(Bach and Harnisch 1979, 6)

I submit that the task of determining adequacy of evidence (and hence of determining the standards imposed by the maxim of Quality) is itself a special case of the sort of phenomenon of which Bach and Harnisch are speaking. In particular, if there is mutual belief that the hearer faces a practical task in which she is in need of information, and that she is relying on the speaker to provide this information, then adequate evidence would be the sort of evidence for a proposition which would render it reasonable for the hearer to act on the assumption that the proposition is true. If there is mutual belief to the effect that the hearer needs information of which she can be certain, then adequate evidence would be the sort of evidence that would support certainty.

The point I am presently making can be formulated in terms of the norm of assertion. Above I presented a schematic version of the hypothesis that assertion has an epistemic norm, in the form of the claim that:

ENA: S must: assert *p*, only if S satisfies epistemic condition E with respect to *p*, i.e. only if S has the *relevant warranting authority* regarding *p*.

My present claim employs the notion of mutual belief to address the matter of what epistemic standards one must satisfy if one is to count as having the relevant warranting authority regarding *p*. To a first approximation (to be modified in section 4 below), the hypothesis is this:

MBN: When it comes to a particular assertion that *p*, the relevant warranting authority regarding *p* depends in part on what is mutually believed by speaker and audience (regarding such things as the participants' interests and informational needs, and the prospects for high-quality information in the domain in question).

Below I will be arguing that (a slightly modified version of) this mutual belief model can address the challenge with which we began this chapter. My only claim here is that the model itself enjoys some independent support: it can be seen as deriving from a broadly Grice-inspired approach to speech exchanges.

This last point is worth dwelling on in a bit more detail. The model I am offering here is a special case of a more general picture of speech exchanges, one having nothing in particular to do with philosophy (or with disagreement, for that matter). On this picture, speech is a cooperative activity, and assertion is to be understood in these terms, as governed by rules of the sort Grice articulated in his principle of Quality. Insofar as these rules are themselves an object of (perhaps merely implicit) mutual belief, they determine a set of mutual expectations of speaker and hearer. That these expectations are (in part) epistemic, demanding adequacy of evidence, is precisely what makes assertion apt for playing the very important role it does: that of serving as the vehicle for the transmission of information. We can bring this out as follows. A hearer who observes a speaker make an assertion, under conditions in which the rules governing assertion are objects of mutual belief, will expect that the speaker acknowledges these rules, and so will expect the speaker to acknowledge the responsibility for having had adequate evidence. Insofar as the hearer regards the speaker as having succeeded at following the rules, then, the hearer regards the speaker as having adequate evidence; and when the hearer's so regarding the speaker is rational, this rationalizes the hearer's move to accept the information the speaker presented in her assertion, on the basis of her having so asserted.⁹ In sum, it is because of the rules governing assertion that this speech act is apt for rationalizing hearers' beliefs in what is asserted—and precisely this renders assertion apt for the transmission of information.

⁹ Of course, if the hearer was irrational in regarding the speaker as having conformed to the rules—the speaker asserted something regarding which it is common knowledge that no one has any evidence, or she had obvious vested interests in getting the hearer to believe what she said, etc.—then the hearer's acceptance is itself rationally flawed.