


The background of the cover features a complex, abstract geometric pattern of black lines on a light gray background. The pattern consists of various rectangles, squares, and triangles, some of which are nested or overlapping, creating a sense of depth and architectural structure. The lines are of varying thicknesses and orientations, contributing to a dynamic and modern aesthetic.

Routledge Studies in Epistemology

THE EPISTEMOLOGY OF GROUP DISAGREEMENT

Edited by
Fernando Broncano-Berrocal
and J. Adam Carter



The Epistemology of Group Disagreement

This book brings together philosophers to investigate the nature and normativity of group disagreement. Debates in the epistemology of disagreement have mainly been concerned with idealized cases of peer disagreement between individuals. However, most real-life disagreements are complex and often take place within and between groups. Ascribing views, beliefs, and judgments to groups is a common phenomenon that is well researched in the literature on the ontology and epistemology of groups. The chapters in this volume seek to connect these literatures and to explore both intra- and inter- group disagreements. They apply their discussions to a range of political, religious, social, and scientific issues. *The Epistemology of Group Disagreement* is an important resource for students and scholars working on social and applied epistemology; disagreement; and topics at the intersection of epistemology, ethics, and politics.

Fernando Broncano-Berrocal is a Talent Attraction Fellow at the Autonomous University of Madrid, Spain. He works mainly in epistemology, with an emphasis on virtue epistemology, philosophy of luck, social epistemology, and collective epistemology. He is the co-editor, with J. Adam Carter, of *The Epistemology of Group Disagreement* (Routledge, 2021). His work has appeared in such places as *Philosophical Studies*, *Analysis*, *Synthese*, and *Erkenntnis*.

J. Adam Carter is Reader in Philosophy at the University of Glasgow, UK. His expertise is mainly in epistemology, with particular focus on virtue epistemology, social epistemology relativism, know-how, epistemic luck, and epistemic defeat. He is the author of *Metaepistemology and Relativism* (2016); co-author of *A Critical Introduction to Knowledge-How* (2018); and co-editor, with Fernando Broncano-Berrocal, of *The Epistemology of Group Disagreement* (Routledge, 2021). His work has appeared in *Noûs*, *Philosophy and Phenomenological Research*, *Philosophical Studies*, *Analysis*, and the *Australasian Journal of Philosophy*.

Routledge Studies in Epistemology

*Edited by Kevin McCain, University of Alabama at Birmingham,
USA and Scott Stapleford, St. Thomas University, Canada*

Well-Founded Belief

New Essays on the Epistemic Basing Relation

Edited by J. Adam Carter and Patrick Bondy

Higher-Order Evidence and Moral Epistemology

Edited by Michael Klenk

Social Epistemology and Relativism

*Edited by Natalie Alana Ashton, Martin Kusch, Robin McKenna and
Katharina Anna Sodoma*

Epistemic Duties

New Arguments, New Angles

Edited by Kevin McCain and Scott Stapleford

The Ethics of Belief and Beyond

Understanding Mental Normativity

Edited by Sebastian Schmidt and Gerhard Ernst

Ethno-Epistemology

New Directions for Global Epistemology

Edited by Masaharu Mizumoto, Jonardon Ganeri, and Cliff Goddard

The Dispositional Architecture of Epistemic Reasons

Hamid Vahid

The Epistemology of Group Disagreement

Edited by Fernando Broncano-Berrocal and J. Adam Carter

For more information about this series, please visit: <https://www.routledge.com/Routledge-Studies-in-Epistemology/book-series/RSIE>

The Epistemology of Group Disagreement

Edited by
Fernando Broncano-Berrocal
and J. Adam Carter

First published 2021
by Routledge
52 Vanderbilt Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2021 Taylor & Francis

The right of Fernando Broncano-Berrocal and J. Adam Carter to be identified as the authors of the editorial material, and of the authors for their individual chapters, has been asserted in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data
A catalog record for this title has been requested

ISBN: 978-0-367-07742-6 (hbk)
ISBN: 978-0-429-02250-0 (ebk)

Typeset in Sabon
by codeMantra

Contents

1 The Epistemology of Group Disagreement: An Introduction	1
FERNANDO BRONCANO-BERROCAL AND J. ADAM CARTER	
2 Deliberation and Group Disagreement	9
FERNANDO BRONCANO-BERROCAL AND J. ADAM CARTER	
3 Disagreement within Rational Collective Agents	46
JAVIER GONZÁLEZ DE PRADO SALAS AND XAVIER DE DONATO-RODRÍGUEZ	
4 When Conciliation Frustrates the Epistemic Priorities of Groups	68
MATTIAS SKIPPER AND ASBJØRN STEGLICH-PETERSEN	
5 Intra-Group Disagreement and Conciliationism	90
NATHAN SHEFF	
6 Bucking the Trend: The Puzzle of Individual Dissent in Contexts of Collective Inquiry	103
SIMON BARKER	
7 Gender, Race, and Group Disagreement	125
MARTIN MIRAGOLI AND MONA SIMION	
8 Disagreement and Epistemic Injustice from a Communal Perspective	139
MIKKEL GERKEN	
9 Group Disagreement in Science	163
KRISTINA ROLIN	

10 Disagreement in a Group: Aggregation, Respect for Evidence, and Synergy	184
ANNA-MARIA ASUNTA EDER	
11 Why Bayesian Agents Polarize	211
ERIK J. OLSSON	
12 The Mirage of Individual Disagreement: Groups Are All that Stand between Humanity and Epistemic Excellence	230
MAURA PRIEST	
13 A Plea for Complexity: The Normative Assessment of Groups' Responses to Testimony	259
NIKOLAJ NOTTELMANN	
<i>Notes on Contributors</i>	281
<i>Index</i>	283

1 The Epistemology of Group Disagreement

An Introduction

*Fernando Broncano-Berrocal and
J. Adam Carter*

1.1 Group Disagreement: A Brief Overview

Disagreement is among the most thriving topics in mainstream and social epistemology.¹

The research question responsible for initially launching the *epistemology of disagreement* as its own subfield in the early 2000s can be put very simply: suppose you believe some proposition, p , is true. You come to find out that an individual whom you thought was equally likely as you are to be right about whether p is true, believes *not- p* . What should you do? Are you rationally required, given this new evidence, to revise your initial belief that p , or is it rationally permissible to simply ‘hold steadfast’ to your belief that p with the same degree of confidence that you did before you found out your believed-to-be epistemic peer disagreed with you? Call this the *peer disagreement question*.

How we go about answering this question has obvious practical ramifications: we disagree with people we think are our peers often; knowing what we should do, epistemically, would be valuable guidance. But the peer disagreement question is also important for epistemologists to understand, theoretically speaking, given that it has direct ramifications for how we should understand disagreement itself as a form of evidence.

Unsurprisingly, responses to the peer disagreement question have fallen into two broadly opposing categories: those who think that discovering that an epistemic peer disagrees with you rationally requires of you some substantial kind of *conciliation*²—perhaps even agnosticism³—and those who think that it does not.⁴ Interestingly, the past ten years or so have shown that—in the close orbit of the peer disagreement question—there are a range of related and interesting epistemological questions, questions that are perhaps just as epistemologically as well as practically significant.⁵

Just consider that the peer disagreement question is *individualistically framed*. It is a question about what rationality requires of an individual when they disagree with another individual about some contested proposition. Gaining an answer tells us, *at most*, and in short, what

individuals should do in the face of epistemic adversity. But we also want to know what *groups* should do in the face of epistemic adversity. For example: what should a group—say, a scientific committee—do if it turns out that one of the members on the committee holds a view that runs contrary to the consensus?⁶

It would be convenient if answering questions about how individuals should respond to epistemic adversity *implied* answers to the interesting questions about how groups should do the same. Unfortunately, though, things are not so simple. This is because, to a first approximation, the epistemic properties of groups are not, as recent collective epistemology has suggested, always simply reducible to an aggregation of the epistemic properties of its members.⁷ If we want to understand what *groups* should do, rationally speaking, when there is internal disagreement among members, or when there is disagreement between a group and individuals or groups external to the group, we cannot and should not expect to find the answers we need simply by looking to the results social epistemology has given us to questions that were *individualistically framed*.

The topic of this volume—the *epistemology of group disagreement*—aims to face the complex topic of group disagreement head on; it represents the first-ever volume of papers dedicated exclusively to *group disagreements* and the epistemological puzzles such disagreements raise. The volume consists of 12 new essays by leading epistemologists working in the area, and it spans a range of different key themes related to group disagreement, some established themes and others entirely new. In what follows, we offer brief summaries of these 12 chapters, drawing some connections between them where appropriate.

1.2 Overview of Chapters

In general, there are two epistemically significant ways in which intragroup disagreement can be resolved, i.e., in which members of a divided group can come to agree to let a certain view stand as the group's view: (i) they can deliberate and/or (ii) take a vote. Which is the best strategy and why? In 'Deliberation and Group Disagreement', we (Fernando Broncano-Berrocal and J. Adam Carter) open the volume by exploring the epistemic significance that the key difference between deliberative and voting procedures has for the resolution of intragroup disagreement: namely, the fact that only deliberation necessarily requires that group members communicate with each other and by doing so exchange their evidence. In order to make traction on this question, deliberation's epistemic effectiveness in resolving intragroup disagreement is assessed in some detail with respect to how well, in comparison with voting, it promotes (or thwarts the attainment of) a range of different epistemic goals, including truth, evidential support, understanding, and epistemic justice.

Javier González de Prado Salas and Xavier de Donato-Rodríguez, in their contribution 'Disagreement Within Rational Collective Agents', are primarily concerned with the question of what a group must do to be rational as a group when members of that group hold disagreeing views. One answer that they consider and reject holds that group attitudes are rational if they result from the application of appropriate judgement aggregation methods. On the proposal they favour, group (epistemic) attitudes are rational insofar as they are formed by responding competently or responsibly to the (epistemic) reasons available to the group as a group, where this requires the exercise of reasons-responding competences attributable to the group. In developing this proposal, González de Prado and de Donato-Rodríguez defend conciliationism as having an important role to play, and offer a positive characterization of group deliberation according to which deliberation in collective agents tends to facilitate the achievement of internal agreement, not only about what attitude to adopt collectively but also about the reasons for doing so.

Whereas González de Prado and de Donato-Rodríguez helpfully show the positive implications of conciliationism about group disagreement—in that it offers an optimistic picture of collective deliberation as a rational method for intragroup disagreement resolution—Mattias Skipper and Asbjørn Steglich-Petersen highlight its shortcomings. In particular, in their chapter 'When Conciliation Frustrates the Epistemic Priorities of Groups', Skipper and Steglich-Petersen argue that conciliatory responses on behalf of individual group members to intragroup disagreement—even if rational *qua* response types to individual disagreement—can have adverse epistemic consequences at the group level. In particular, as they see it, the problem is that such conciliatory responses to an internal disagreement can frustrate a group's epistemic priorities by changing the group's relative degree of reliability in forming true beliefs and avoiding false ones. Finally, Skipper and Steglich-Petersen suggest a solution to this epistemic priority problem that does not imply abandoning conciliationism.

The next two papers in the volume continue follow suit in investigating the relationship between group disagreement and conciliationism, albeit in different ways. In his chapter 'Intra-Group Disagreement and Conciliationism', Nathan Sheff's objective is to defend a form of conciliationism in the specific context of intra-group disagreements. Conciliationism in this context holds when an individual dissenter finds herself in disagreement with the other members of a deliberative group, the rational response for the disagreeing member is lowering confidence in their view. Sheff argues first that (i) intra-group conciliationism does not enjoy *ex ante* the intuitive plausibility that ordinary conciliationism, viz., individualistically framed, does, but (ii) difficulties facing the view can be overcome when we suitably appreciate, with reference to Margaret Gilbert's account of joint commitment,⁸ the kind of normativity that constrains an individual dissenter in the predicament of an intragroup disagreement. In particular,

they find themselves epistemically responsible for contradictory views: their own view, and that of the group and accordingly pulled in contrary directions. In this circumstance, Sheff argues, the rational response is at least to lower their confidence in their view.

In 'Bucking the Trend: The Puzzle of Individual Dissent in Contexts of Collective Inquiry', Simon Barker, like Sheff, is concerned with the predicament of an individual dissenter in her capacity as a group member. As Barker observes, there is pressure to suppose that when an individual dissents with intragroup members, the greater the number of one's peers against one, the more significance one should afford the disagreement—viz., what he calls the principle of collective superiority. At the same time, he notes, discussions of disagreement within collective inquiry have maintained that justified collective judgements demand methods of inquiry that permit and preserve (rather than eliminate) dissent—viz., a principle that Barker labels epistemic liberalism. Taken together, these principles seem to make different and incompatible demands, what Barker calls the 'puzzle of individual dissent'. Barker's objective in the paper is to sharpen this puzzle by tracing out the consequences of rejecting either of the two principles jointly responsible for the dilemma, and to assess the significance of the dilemma more widely in epistemology.

The next three papers in the volume engage in different ways with the social and power dynamics of group disagreement. In 'Gender, Race, and Group Disagreements' Mona Simion and Martin Miragoli take as a starting point two cases of group disagreement, one involving gender discrimination, the other involving the marginalization of racial and religious minorities. Both, they argue, feature a distinctive form of epistemic injustice at play, and further, that extant views in the epistemology of peer disagreement have difficulties accounting for what is defective about these cases. Against this background, Simion and Miragoli propose and defend a two-tiered solution to the problem that relies on an externalist epistemology and a functionalist theoretical framework.

Epistemic injustice is also a central theme in Mikkel Gerken's contribution to the volume, 'Disagreement and Epistemic Injustice from a Communal Perspective'. Gerken's central focus is on the epistemic pros and cons of disagreement for a community and on how the social structure of the community bears on these pros and cons. A central conclusion drawn is that disagreement has more epistemic costs at the communal level than is often recognized by those who follow Mill's emphasis on disagreement's positive social significance, and that these epistemic costs often yield epistemic injustice, especially given disagreement's capacity to defeat testimonial warrant.

In 'Group Disagreement in Science', Kristina Rolin explores, through the lens of scientific dissent, how relations of power influence perceived epistemic responsibilities. Rolin takes as a starting point the widespread view in the philosophy of science that a scientific community has an obligation

to engage scientific dissent only when it is normatively appropriate from an epistemic point of view. One notable line of criticism to this standard line maintains that the norms constraining epistemically appropriate dissent are ambiguous. Rolin's objective is to respond to this concern by arguing that even when there is disagreement over the interpretation of such norms, a scientific community has a moral reason to respond to dissenters. On her favoured approach, there is a norm of epistemic responsibility—both an epistemic and moral norm—that defines mutual obligations for dissenters and the advocates of a consensus view.

The volume's next two chapters view the epistemology of group disagreement through a more formal lens. In 'Disagreement in a Group: Aggregation, Respect for Evidence, and Synergy', Anna-Maria Asunta Eder seeks to answer the following guiding question: How do members of a group reach a rational epistemic compromise on a proposition when they have different rational credences in the proposition? One way to settle this question is a standard Bayesian method of aggregation, a commitment of which is that the only factors among the agents' epistemic states that matter for finding the compromise are the group members' credences. In contrast, Eder develops and defends a different approach—one that makes use of a fine-grained method of aggregation—on which the members' rational credences are not the only factors concerning the group agents' rational epistemic states that matter for finding an epistemic compromise. This method is based on a non-standard framework for representing rational epistemic states that is more fine-grained than Standard Bayesianism, and which comports with a Dyadic Bayesian framework Eder has defended in a previous work.⁹

A different kind of Bayesian approach to group disagreement is explored by Erik J. Olsson in his paper 'Why Bayesian Agents Polarize'. A number of studies have concluded that ideal Bayesian agents can end up seriously divided on an issue given exactly the same evidence, which suggests that polarization may be rational. But even if this is right, a separate question is *why* do Bayesian agents polarize? Olsson engages with this question in the context of the Bayesian Laputa model of social network deliberation, developed by Angere and Olsson (e.g., 2017). According to recent work by Pallavicini, Hallsson, and Kappel (2018), on the Laputa model, polarization arises due to a failure of Laputa to take into account higher-order information in a particular way, making the model incapable of capturing full rationality. Olsson's objective is to reject Pallavicini et al.'s argument; on his preferred assessment, what drives polarization is expectation-based updating in combination with a modelling of trust in a source that recognizes the possibility that the source is systematically biased.

The volume rounds out with two new spins on traditional ways of thinking about groups and evidence in cases of (group) disagreement. In her paper 'The Mirage of Individual Disagreement: Groups Are All

that Stand between Humanity and Epistemic Excellence’, Maura Priest argues that a large number of important and long-standing disagreements that have typically been understood as between individuals, are actually disagreements between collectives. This conclusion marks a departure from orthodox thinking about individual disagreement. But, once this is appreciated, she argues, it is easier to then appreciate why such disagreements are often long-standing; further, Priest argues, many individual disagreements (properly understood as group disagreements) are likely to remain unresolved because the relevant parties are not properly motivated by epistemic ends.

The volume ends with Nikolaj Nottelmann’s paper, ‘A Plea for Complexity: The Normative Assessment of Groups’ Responses to Testimony’. Nottelmann’s central aim is to show that the epistemic evaluation of group performance in the face of testimony and disagreement is a more complex matter than has so far been explicitly acknowledged in the literature. In many cases, he argues, it is far from clear whether our evaluations of a group’s responses to testimony are primarily epistemic or moral, and, in the latter case, how epistemic standards play into our moral assessment. In addition, Nottelmann maintains that what count as the relevant criteria of groupness, group membership, and group belief vary according to our evaluative interests and perspectives.¹⁰

Notes

- 1 For some representative recent work on disagreement in epistemology, see, for example, Carey and Matheson (2013); Christensen (2007, 2009); Elga (2007); Feldman (2007); Feldman and Warfield (2010); Goldman (2010); Hales (2014); Kelly (2005); Lackey (2013); Littlejohn (2013); MacFarlane (2007); Matheson (2009, 2015, 2016); Sosa (2011); Thune (2010a, 2010b); and Carter (2018).
- 2 This view is often described as ‘conciliationism’. See, e.g., Feldman (2007) and Elga (2007).
- 3 See Feldman (2007); cf., Carter (2018).
- 4 For some representative ‘non-conciliationist’ views, see, e.g., Kelly (2005); Foley (2001), and Wedgwood (2007).
- 5 One notable example here concerns the *uniqueness thesis* (e.g., Kelly 2013; Dogramaci and Horowitz 2016; Matheson 2011) which holds that, with respect to a proposition *p*, your body of evidence, *E*, justifies at most one of the three attitudes of belief, disbelief, and withholding vis-a-vis *p*. For criticism of uniqueness, see, e.g., Kelly (2005); Ballantyne and Coffman (2011), and Goldman (2010).
- 6 For discussion on this issue, see, e.g., the essays in Lackey (2014) and Brady and Fricker (2016).
- 7 See, e.g., Gilbert (1996, 2013); Tollefsen (2006, 2007, 2015), Tuomela (1995, 2002, 2013), and Palermos (2015).
- 8 See, e.g., Gilbert (1996, 2013).
- 9 Brössell and Eder (2014).
- 10 Fernando Broncano-Berrocal is grateful to the BBVA Foundation for supporting this book, which was edited as part of a 2019 BBVA Leonardo Grant for Researchers and Cultural Creators—the BBVA Foundation accepts no

responsibility for the opinions, statements, and contents included in this book, which are entirely the authors' responsibility. J. Adam Carter is grateful to the Leverhulme Trust for supporting this book, which was edited as part of the Leverhulme-funded 'A Virtue Epistemology of Trust' (RPG-2019-302) project, hosted by the University of Glasgow's COGITO Epistemology Research Centre.

References

- Angere, Staffan, and Erik J. Olsson. 2017. 'Publish Late, Publish Rarely!: Network Density and Group Performance in Scientific Communication'. In *Scientific Collaboration and Collective Knowledge*, edited by T. Boyer-Kassem, C. Mayo-Wilson, and M. Weisberg, 34–62. Oxford: Oxford University Press.
- Ballantyne, Nathan, and E. J. Coffman. 2011. 'Uniqueness, Evidence, and Rationality'. *Philosophers' Imprint* 11 (8): 1–13.
- Brady, Michael S., and Miranda Fricker. 2016. *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*. Oxford: Oxford University Press.
- Brüssel, Peter, and Anna-Maria A. Eder. 2014. 'How to Resolve Doxastic Disagreement'. *Synthese* 191 (11): 2359–2381.
- Carey, Brandon, and Jonathan Matheson. 2013. 'How Skeptical Is the Equal Weight View?: Brandon Carey and Jonathan Matheson'. In *Disagreement and Skepticism*, edited by Diego E. Machuca, 140–158. London: Routledge.
- Carter, J. Adam. 2018. 'On Behalf of Controversial View Agnosticism'. *European Journal of Philosophy* 26 (4): 1358–1370.
- Christensen, David. 2007. 'Epistemology of Disagreement: The Good News'. *Philosophical Review* 116 (2): 187–217.
- . 2009. 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass* 4 (5): 756–767.
- Dogramaci, Sinan, and Sophie Horowitz. 2016. 'An Argument for Uniqueness about Evidential Support'. *Philosophical Issues* 26 (1): 130–147.
- Elga, Adam. 2007. 'Reflection and Disagreement'. *Noûs* 41 (3): 478–502.
- Feldman, Richard. 2007. 'Reasonable Religious Disagreements'. In *Philosophers without Gods: Meditations on Atheism and the Secular*, edited by Louise Anthony, 194–214. Oxford: Oxford University Press.
- Feldman, Richard, and Ted A. Warfield. 2010. *Disagreement*. Oxford: Oxford University Press.
- Foley, Richard. 2001. *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Gilbert, Margaret. 1996. *Living Together: Rationality, Sociality, and Obligation*. London: Rowman & Littlefield Publishers.
- . 2013. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.
- Goldman, Alvin I. 2010. 'Epistemic Relativism and Reasonable Disagreement'. In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 187–215. Oxford: Oxford University Press.
- Hales, Steven D. 2014. 'Motivations for Relativism as a Solution to Disagreements'. *Philosophy* 89 (1): 63–82.
- Kelly, Thomas. 2005. 'The Epistemic Significance of Disagreement'. *Oxford Studies in Epistemology* 1: 167–196.

- Kelly, Thomas. 2013. 'Evidence Can Be Permissive'. In *Contemporary Debates in Epistemology*, edited by M. Steup and J. Turri, 298. Blackwell.
- Lackey, Jennifer. 2014. *Essays in Collective Epistemology*. Oxford: Oxford University Press.
- Lackey, Jennifer Amy. 2013. 'Disagreement and Belief Dependence: Why Numbers Matter'. In *The Epistemology of Disagreement: New Essays*, edited by D. Christensen and J. Lackey, 243–268. Oxford: Oxford University Press.
- Littlejohn, Clayton. 2013. 'Disagreement and Defeat'. In *Disagreement and Skepticism*, edited by D. Machuca, 169–193. Routledge.
- MacFarlane, John. 2007. 'Relativism and Disagreement'. *Philosophical Studies* 132 (1): 17–31.
- Matheson, Jonathan. 2009. 'Conciliatory Views of Disagreement and Higher-Order Evidence'. *Episteme* 6 (3): 269–279.
- . 2011. 'The case for rational uniqueness'. *Logos & Episteme* 2 (3): 359–373.
- . 2015. *The Epistemic Significance of Disagreement*. Springer.
- . 2016. 'Moral Caution and the Epistemology of Disagreement'. *Journal of Social Philosophy* 47 (2): 120–141.
- Pallavicini, Josefine, Björn Hallsson, and Klemens Kappel. 2018. Polarization in groups of Bayesian agents. *Synthese*. <https://doi.org/10.1007/s11229-018-01978-w>.
- Palermos, Spyridon Orestis. 2015. 'Active Externalism, Virtue Reliabilism and Scientific Knowledge'. *Synthese* 192 (9): 2955–2986. <https://doi.org/10.1007/s11229-015-0695-3>.
- Sosa, Ernest. 2011. 'The Epistemology of Disagreement'. In *Social Epistemology*, edited by A. Haddock, A. Millar and D. Pritchard. Oxford University Press.
- Thune, Michael. 2010a. "Partial Defeaters" and the Epistemology of Disagreement'. *The Philosophical Quarterly* 60 (239): 355–372.
- . 2010b. 'Religious Belief and the Epistemology of Disagreement'. *Philosophy Compass* 5 (8): 712–724.
- Tollefsen, Deborah. 2007. 'Group Testimony'. *Social Epistemology* 21 (3): 299–311.
- Tollefsen, Deborah Perron. 2006. 'From Extended Mind to Collective Mind'. *Cognitive Systems Research* 7 (2–3): 140–150.
- . 2015. *Groups as Agents*. Hoboken, NJ: John Wiley & Sons.
- Tuomela, Raimo. 1995. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press.
- . 2002. *The Philosophy of Social Practices: A Collective Acceptance View*. Cambridge University Press.
- . 2013. *Cooperation: A Philosophical Study*. Vol. 82. Dordrecht: Springer Science & Business Media.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.

2 Deliberation and Group Disagreement

*Fernando Broncano-Berrocal
and J. Adam Carter*

2.1 Setting the Stage: Deliberative versus Non-Deliberative Agreement Following Intragroup Disagreement

Many disagreements take place in group settings. Over the years, religious groups (e.g., Christians) have internally disputed topics they consider significant (e.g., the real presence of Christ in the Eucharist). More often than not, political parties (e.g., the Tories) go through internal divisions over issues of societal importance (e.g., a no-deal Brexit). A brief look at the history of science reveals how scientists (e.g., physicists) disagree over factual issues in their fields (e.g., the Copenhagen vs. the many-worlds interpretation of quantum mechanics). More mundanely, disputes over practical matters are the order of the day in many families. On occasions, such internal disagreements end up badly, with a split in the relevant group or a punishment for the less influential. Sometimes, however, they result in a consensus¹ or an agreement of sorts to take a particular course of action or to let some view stand as the group's view.² It is this latter kind of intragroup disagreement we are interested in: the one that gets resolved.

How members of a group internally disagree matters for many reasons, not only for the stability or survival of the group but also *epistemically*. In general, there are two epistemically significant ways in which intragroup disagreement can be resolved, i.e., in which members of a divided group can come to *agree to let a certain view stand as the group's view*: (i) they can *deliberate* and/or (ii) take a *vote*.

In this chapter, we are interested in investigating the epistemic significance that the key difference between deliberative and voting procedures has for the resolution of intragroup disagreement: namely, the fact that only deliberation necessarily requires that group members communicate with each other and, more specifically, the fact that, by doing so, they exchange their evidence. Thus, the paper aims to assess, in general, the epistemic significance that such an exchange (or lack thereof) has for the resolution of intragroup disagreement.

This is of course not to say that deliberation and voting are mutually exclusive mechanisms for groups to resolve their internal disputes.

In practice, groups settle their disagreements by *mixed methods of decision-making*, i.e., methods that both involve deliberating and voting—as is, for instance, the mixed method for judging articles of impeachment in the United States House of Representatives.

That said, to better pin down the epistemic significance of each, it is best to keep them apart, at least theoretically. Thus, the kind of cases we will mainly focus on (whether real or ideal) have the following structures:

Deliberative cases of intragroup disagreement: Some operative members³ of group *G* hold *p* and some not-*p* at *t*₁; at *t*₂, *G*'s operative members deliberate among themselves (i.e., they exchange reasons, evidence, arguments, and so on) with an eye toward settling whether *p* or else not-*p* should stand as *G*'s view; at *t*₃, as a result of this process, they settle on either *p* or not-*p*.

Non-deliberative cases of intragroup disagreement: Some operative members of *G* hold *p* and some not-*p* at *t*₁; at *t*₂, *G*'s operative members aggregate their views by taking a vote given some voting rule (e.g., majority rule), absent any communication among each other, with an eye toward settling whether *p* or else not-*p* should stand as *G*'s view; at *t*₃, as a result of this process, they settle on either *p* or not-*p*.⁴

Some clarifications are in order. First, at *t*₂, both in the case of deliberation and voting, members who initially believed that *p* should stand as *G*'s view *may change their opinion*, and vice versa. Second, we leave unspecified the *number of group members* that respectively hold *p* and not-*p* to make it compatible with several possibilities—as we will see, this factor marks a distinction in terms of reliability between deliberation and voting. Third, lack of communication among group members in the case of voting is compatible with there being common knowledge (perhaps implicit) of the existence of an internal disagreement or of the fact that it is to be solved by taking a vote.

Finally, certain cases will not be our main focus. Quite often, members of a group settle on a collective view pursuing *non-epistemic goals*—regardless of whether this collective agreement is reached by deliberation or vote. For example, the board of directors of a pharmaceutical company might settle on the view that a newly marketed drug is not the cause of the death of many, even if they know it, to prevent huge financial losses. A government might systematically deny that the country's secret services have been used for morally contentious surveillance activities, even if known to be true, to prevent protests and media pressure. A religious organization might conceal criminal activities by its members—and thus uphold the collective view that such activities never happened—to avoid criminal charges and loss of reputation.

The reason we won't focus on such cases is that deliberative and voting procedures have little or no epistemic value when aimed at non-epistemic goals.⁵ Instead, the kind of cases we are interested in are cases

of intragroup disagreement in which members of a group reach a collective agreement pursuing *epistemic goals*.⁶ Although this certainly reduces the scope of our inquiry, by idealizing our focus in this way, we will be in a better position to rule out pragmatic noise when answering the two key *epistemological* questions of the paper.⁷ For ease of reference, call these the *resolution question* and the *deliberation question*.

Resolution question: What is the most epistemically appropriate way to resolve intragroup disagreement: by means of deliberation or by taking a vote? More specifically, to what extent is it epistemically advantageous and disadvantageous that group members *exchange evidence* when it comes to reaching a collective agreement?

Deliberation question: Which conditions should *deliberative* disagreement comply with to be epistemically appropriate? More specifically, what would it take to overcome, or at least mitigate, the epistemic disadvantages of resolving intragroup disagreement by means of deliberation?

Our methodological approach to answering these questions is based on a simple *working assumption*: a group's collective endeavor to solve an internal dispute can be aimed at different (albeit not necessarily incompatible) epistemic goals. More carefully:

Assumption: Possibly, for two epistemic goals, E and E^* , and for two groups, G and G^* , members of G would let p or else not- p stand as the G 's view only if the collectively accepted view has epistemic property E and members of G^* would let p or else not- p stand as G^* 's view only if the collectively accepted view has epistemic property E^* .

With this assumption in place, each epistemic goal can be interpreted as providing a particular *standard* for assessing the epistemic significance of deliberating and voting in the resolution of intragroup disagreement. More specifically, the way we propose to assess this epistemic significance is in terms of *goal-conduciveness*: for each goal, we can assess to what extent the fact that group members exchange (or refrain from doing so) reasons and evidence are conducive to it.

The following are four salient candidate epistemic goals we will consider, albeit they are not exhaustive (see n. 9). For any group G in which some operative members hold p and some hold not- p , in trying to settle whether p or else not- p should stand as G 's view by means of method M , members of G would let p or else not- p stand as the group's view only if:

- 1 *Truth*: the collectively accepted view is true.
- 2 *Evidence*: the collectively accepted view is better supported by the best evidence individually possessed by group members than the opposite view.

- 3 *Understanding*: the collectively accepted view leads to more understanding than the opposite view.
- 4 *Epistemic justice*: the fact that *G*'s members let such a view stand as *G*'s view does not wrong any member specifically in her capacity as an epistemic subject (e.g., as a giver of knowledge, in her capacity for social understanding, and so on) or any other person outside the group in that capacity.^{8,9}

Before assessing each epistemic goal, a final methodological caveat is in order. Our approach to the epistemic significance of deliberation and voting in terms of goal-conduciveness does not entail that the different epistemic goals are incompatible with each other, nor does it imply any stance on a number of debates, including (i) whether the relevant goals are finally or instrumentally valuable (or fundamentally or derivatively) in the case of collectives (cf., Goldman 1999; Fallis & Mathiesen 2013); or (ii) whether deliberation has a constitutive epistemic aim in terms of one of these goals; (iii) or whether deliberation has procedural in addition to instrumental epistemic value (cf., Peter 2013). Our results might of course be relevant to these debates, but we stay neutral on them.¹⁰

Here is the plan. In §2.2, we address the truth goal, explain what the different kinds of evidence involved in deliberation are and how they bear on the individual reliability of deliberators; compare the collective reliabilities of deliberation and voting drawing on social choice theory, and show how complex it is to give a straightforward answer to the question of whether deliberation is reliable due to, among other things, the existence of several reliability-undermining group phenomena—which are widely investigated in social psychology. In §2.3, we explain why it shouldn't be assumed that deliberation always achieves optimal results, nor that voting always produces suboptimal outcomes *vis-à-vis* the goal of evidence. In §2.4, we offer two interpretations of the understanding goal and argue that, on both interpretations, deliberation outperforms mere voting. In §2.5, we argue that voting is more efficacious than deliberation with respect to the goal of epistemic justice. In §2.6, we propose several ways to mitigate the potential epistemic disadvantages of solving intragroup disagreement by means of deliberation in relation to each epistemic goal.

2.2 Assessing for Truth

As we have seen, in settling for a collective view, groups may pursue non-epistemic goals (e.g., preventing financial losses), but they sometimes pursue epistemic goals. One example of an epistemically respectable goal, if not the most fundamental epistemic goal,¹¹ is *truth*.

In scientific disagreement, for example, members of a research group that internally disagree over some factual issue would not let a view

stand as the group's view unless they considered it true, or at least, more likely to be true than any competing view. In a quiz show, members of a divided team would not let an answer stand as the team's answer unless they considered it correct (or likely to be correct). Thus, for any given method that members of a divided group may use to reach a collective agreement, we can assess its epistemic propriety in terms of how conducive toward truth this method is. Crucially, the *reliability* of deliberative and voting methods depends on the kind of individual and collective conditions under which they are employed. Fortunately, these conditions have been widely investigated in disciplines such as social psychology and social choice theory. That being so, we will review some of their results (with an eye on truth as the relevant epistemic standard) so as to provide an answer to the resolution question on a safe theoretical and empirical ground.

Before that, it is worth pointing out an *epistemic difference* concerning the reliability of deliberative and voting procedures in general. This will allow us to subsume some relevant results of the aforementioned disciplines under a broader epistemological framework. Consider, first, the following general idea: the reliability of a group in letting only a view that is true stand as the collective view is to some extent premised upon the reliability of individual group members in choosing the right view *both* in the case of deliberation and voting.

To see this, consider a group (e.g. a flat Earth society) such that all members are utterly unreliable (e.g., almost always, they get things wrong) regarding the question of whether p (e.g. whether Earth is flat or spherical). Suppose that this kind of group internally disagrees on whether p or else not- p should stand as the group's view. Even if all members aim at settling on the true view, most will end up defending the false view because of their utter unreliability, whereas those who unlikely end up upholding the right view, will do it by luck. In such a situation, it doesn't matter whether the group deliberates or takes a vote: whatever the procedure for settling their internal disagreement, it will be an unreliable one. Therefore, group members need to be individually reliable to a minimum degree for them to reliably reach a correct collective agreement as a group—in subsection 2.2.2, we will see what the minimum required degree of individual reliability is according to social choice theory.

Nonetheless, while the reliability of a group in letting only a view that is true stand as the collective view is to some extent premised upon the reliability of individual group members both in the case of deliberation and voting, individual reliability is fixed differently in deliberative and non-deliberative cases. The reason, as we will argue next, has to do with the fact that the former involve *different kinds of evidence* besides private evidence and hence different individual competences are required to evaluate them (this will also have a bearing on our discussion in §2.3).

2.2.1 *Individual Reliability*

In non-deliberative cases (at least as we have conceived them), the only evidence that group members use to establish which of the two options in a given dispute (p and not- p) is true (and therefore which one is the one that the group should uphold) is their own *private evidence*. In deliberative cases, by contrast, group members possess not only private evidence, but they are also exposed to *shared evidence*—i.e. evidence bearing on p /not- p shared by other group members during deliberation. In addition, as a consequence of this sharing process, they are also exposed to evidence about the distribution of opinions within the group, or *social evidence* for short—i.e. evidence that n number of group members are in agreement and m number in disagreement with one.

Thus, one plausible idea is that, for any given group member, her overall individual reliability concerning the disputed matter will be determined by how reliable she is in accurately judging to what extent *each* kind of evidence supports p or not- p . Interestingly, precisely because these are different kinds of evidence, the degree of reliability in assessing them need not coincide, hence the divergence with non-deliberative cases. Let's consider each in turn.

First, group members can be more or less reliable at seeing how relevant to the disputed matter their *private evidence* is, and on how much it supports or counts against the views in conflict. If, for instance, a group member's private evidence is misleading evidence for p because, say, it comes from a seemingly trustworthy but ultimately unreliable source, she will hardly assess in a reliable fashion that her evidence does not actually count in favor of p . Suppose that voting is the relevant procedure for resolving intragroup disagreement in a given case. In the difficult cases of misleading evidence just mentioned—as well as in cases where a group member has no evidence whatsoever—not voting for either option might be the best action to avoid collective error.

When it comes to *shared evidence*, matters are more complex. When group members put their private evidence on the table during deliberation, all involved members are exposed to two different things: (i) information that may back up, conflict, be redundant with or even irrelevant to their private evidence, and (ii) judgments from other group members to the effect that the shared information supports p or else not- p to such-and-such degree. Accordingly, and as in the case of private evidence, a group member can be more or less reliable at assessing to what extent the information provided by other members is relevant and corroborative of p or of not- p , and this can be done by, among other things, correctly assessing to what extent those other members are *competent* information-gatherers. Interestingly, as some have noted (e.g. Elga 2010; Weatherson 2013; Eder [this volume]), being competent at acquiring evidence is independent from being competent

at correctly judging the confirmational import of the evidence. Thus, on top of being more or less reliable at assessing the evidence shared during deliberation, group members can be more or less reliable at assessing to what extent their fellow members' assessments of such shared evidence are accurate.

The last kind of evidence involved in deliberative cases, *social evidence*, is somehow different, as it does not directly bear on the question of whether p . As we have defined it, social evidence is evidence about, specifically, the distribution of opinions within the group, i.e., evidence about how many group members believe that p (or else not- p) should stand as the group's view because p (or else not- p) is true. Interestingly, social evidence can have a defeating effect on its own even if it does not directly bear on the question of whether p , and namely even if it carries no more information than that of assertions of the type "I think that p should stand as the group's view".

To see this, suppose that you have conclusive private evidence for the truth of p and that on that basis you believe that p should stand as the group's view. Furthermore, suppose that you are the only person in your group in possession of evidence that is relevant to the disputed matter. You share your evidence with your fellow members, which you regard as your epistemic peers. Suppose, next, that no one is moved and all of them (e.g., 999 members), except for you, individually assert "I believe that not- p should stand as the group's view because not- p is true". Many in the epistemological literature on disagreement agree that you should reduce your confidence in your belief simply because of being exposed to social evidence to the effect that a majority is in disagreement with you.¹² Furthermore, this defeating effect may occur even if private evidence directly bearing on p /not- p has not been put on the table yet.¹³

Turning to reliability, the kind of competence required to judge whether the social evidence available in the group is misleading or on the right track is a competence to judge whether the other group members are being *sincere* in asserting things such as "I believe that p should stand as the group's view". Suppose that, during deliberation, someone in your group asserts that. The questions you should ask yourself, *qua* group member, are like these: Is this person being sincere? Does she really care about the truth? Or is she making that assertion for strategic or pragmatic reasons? If one can answer these kinds of questions correctly for all (or at least many) group members, one is reliable at processing the group's available social evidence.¹⁴ By contrast, if one conciliates with the majority for a non-epistemic reason such as social comparison (e.g., to maintain a socially favorable position within the group), one is not reliable at processing the group's social evidence.¹⁵

Which of these different bodies of evidence (i.e., private, shared, or social) and which of the corresponding degrees of reliability in processing them should have a greater weight in the overall individual reliability

of a given group member is a question whose answer hangs to a great extent on the correctness of the different views of the epistemology of disagreement. In general, *steadfast theorists* will be more inclined to assign a greater weight to the group members' assessments of their own private evidence (or even of the shared evidence), whereas *conciliationists* will lean in the direction of giving a greater significance to the judgments of other group members and to the distribution of opinions within the group.¹⁶

To summarize the discussion so far, we've seen that there is a significant difference between how individual reliability is fixed in deliberative and non-deliberative cases of intragroup disagreement. This difference has to do with the fact that, in non-deliberative cases, group members only need to evaluate the confirmational import of their own private evidence to choose between the true and the false collective view. In deliberative cases, by contrast, group members need to process, besides their private evidence, the evidence shared by others as well as the available evidence about the distribution of opinions within the group, which can have a defeating effect by itself. In such deliberative cases, however, it is an open question which kind of evidence should carry more weight in fixing the individual reliability of group members, or what the interplay between these three types of evidence might be.

2.2.2 *Group Reliability*

One question we can ask concerning the reliability of deliberative and voting procedures is how reliable individual members of a group undergoing an internal dispute need be in order for one such procedure to reliably lead the group to settle on the true view. The results of social choice theory become useful on this score.

Let's consider voting first. In general, political scientists assess voting rules in terms of *fairness criteria*, i.e., how sensitive they are to all of the voter's opinions and preferences in the right way (Pacuit 2019). However, interestingly for our epistemological purposes, they can also be assessed in terms of how well they track the truth, i.e., in terms of how much the resulting collective view approximates it (List 2013; Pacuit 2019). A voting rule that is often referred to as a *collective truth-tracking device* in the two-option case (the one we are concerned with) is *majority rule* (see, e.g., List & Goodin 2001).

As it is well-known, one prominent argument for adopting majority rule comes from the Condorcet Jury Theorem (CJT), which maintains that, given two possible positions p and not- p with respect to a given topic (e.g., a verdict, a diagnosis, a factual issue), where only one of the options is correct given some standard (in our case, truth), the probability that a majority votes for the correct option increases and converges to one as the size of the group grows. Crucially, CJT is premised upon

two conditions: (i) that the probability (viz. reliability) that each group member identifies the correct position is greater than 0.5 and the same for all voters (*voter competence condition*) and (ii) that all correct votes are mutually independent, conditional on the truth,¹⁷ which is either p or not- p (*voter independence condition*).¹⁸

Whether or not majority rule reliably yields epistemically appropriate results (true or accurate group agreements) crucially depends on the voter competence and independence conditions being met. But this seldom happens. For instance, factors that have been cited as leading to correlated votes include opinion leaders, schools of thought, communication among voters, or common information (cf., e.g., Ladha 1992). Moreover, as Dietrich and Spiekermann (2020) point out, *any* common cause of votes is a potential source of dependence, including non-evidential (e.g., situational) factors such as distracting heat.¹⁹

Lack of independence has implications not only for whether or not CJT applies to a group that aims to resolve an internal dispute by taking a vote according to majority rule, but also for how the nature of such a disagreement should be conceived. After all, if the votes of every member in the two disagreeing subgroups are correlated, intragroup disagreement comes down to a one-to-one disagreement situation, as there would be two sets of mutually dependent votes: those for p and those for not- p . This is epistemically significant. For one epistemic benefit of CJT is that the *larger* the group, the *better* at tracking the truth it is. Therefore, if all votes are correlated in the two disagreeing subgroups, the size of the group no longer has a bearing on its reliability.

Turning to voter competence, multiple specific factors can bear on the individual reliability of voters. From a general epistemological point of view, the quality of their private evidence is probably the most significant factor. But note that voter reliability can be low even when the evidence privately possessed is good evidence. For the probability that a voter *correctly* judges that her good evidence is supportive of p rather than not- p is independent of the epistemic goodness of the evidence (e.g., someone with conclusive private evidence might fail to notice that the evidence is conclusive because of not being sufficiently attentive). Conversely, a voter with misleading evidence might uncritically follow her evidence, thus making it unlikely that she votes for the correct view. Finally, voters can also be unlikely to vote for the right view when they possess no evidence whatsoever (e.g., by casting their votes for p or not- p merely on the basis of the results of tossing coins that unbeknownst to them are independently biased in favor of the false view). Interestingly, all this can happen while all group members vote with the aim of choosing the correct view.

Lack of voter competence bears on the epistemic appropriateness of majority vote as a way of solving a group's internal dispute. One thing that the literature on CJT shows is that when the votes are independent

but the competence of all voters is lower than 0.5 and the same for all, or when the average judgmental competence of voters in the group is lower than 0.5 (cf., e.g., Grofman et al. 1983), the probability that a majority votes for the correct option increases and converges to 0 as the size of the group grows. So majority rule can be an epistemically *inappropriate* procedure for solving intragroup disagreement after all.

Of course, the literature is filled with jury theorems that relax the independence (Ladha 1992, Dietrich & List 2004; Dietrich & Spiekermann 2013; Goodin & Spiekermann 2018) and the competence (e.g. Grofman et al. 1983; Boland 1989) conditions while still serving as truth-tracking devices, and hence as epistemically appropriate ways to resolve intragroup disagreement, at least in the case of large groups.²⁰ But, in general, voters need to be individually reliable to a sufficient degree, where in most cases this means being better than random.²¹

For some internal disputes, when individual reliability is an issue, groups can opt for some sort of proxy voting system that allows delegation of the votes to the most competent or well-informed in the group, or to weighted majority rules (e.g., expert rules) that assign different weights to different competence distributions (e.g., more weight to the votes of the most competent members). In general, for any competence distribution, there will be an optimal voting truth-tracking procedure for the group to solve its internal disagreement (for optimal voting rules see, e.g., Nitzan & Paroush 1982; Gradstein & Nitzan 1986; Dietrich 2006).

Beyond specific voting rules groups might use to resolve their internal disputes, one question we can ask is this: does communication and evidence sharing among group members represent a significant epistemic advantage over members simply taking a vote on the basis of her private evidence? What are the epistemic benefits and drawbacks of deliberation in general *vis-à-vis* the goal of reaching a true collective agreement?

One way to answer these questions is to offer a formal analysis of deliberation and compare it with voting procedures. Hartmann and Rafée Rad (2018) do precisely this and show that deliberation is truth-conducive in a similar way as majority voting as per CJT. It is worth considering their proposed Bayesian model of deliberation, not only because its results are relevant to the subject matter, but also because it will serve to illustrate the many complexities that communication among group members may give rise to and, therefore, that any formal model of deliberation might need to incorporate if deliberation is to be compared to voting in a realistic way.

Hartmann and Rafée Rad's Bayesian model of deliberation is based on several assumptions. First, all evidence is put on the table *before* deliberation (i.e., no extra evidence shows up during deliberation, so all evidence is shared evidence). Second, group members are assigned a first-order reliability value that measures how correctly they judge the disputed matter. Third, they are assigned a second-order reliability value that reflects how well they estimate the first-order reliability of the other group members. The latter is kept fixed during the course of

deliberation, while the former may increase as members learn to better judge the reliability of other group members. In this way, deliberation, as they model it, consists in the following process:

The group has to decide on the truth or falsity of a hypothesis H. Each group member assigns a certain probability to H. Then each group member casts a vote on the basis of this probability. Then each group member updates her probability on the basis of the votes of the other group members, weighted according to the estimated reliabilities (...). The procedure is iterated, and in each round the second order reliabilities are increased which leads to a more accurate estimation of the reliability of the votes of the other group members. After a number of rounds, this process converges.

(Hartmann & Rafiee Rad 2018: 1278)

Their results show that the truth-tracking properties of deliberation are very similar to those of majority vote. As they summarize them:

The deliberation process results in a consensus and correctly tracks the truth for groups of large size in the following cases: (i) homogeneous groups with a first order reliability greater than 0.5 and with a high second order reliability (ii) inhomogeneous groups with average first order reliabilities above 0.5 and with a high (initial) second order reliability. In this sense the deliberation procedure manifests the same epistemic properties as the majority voting while adding the benefit of a group consensus (...) We furthermore provided simulation results that indicate that the deliberation procedure tracks the truth even in cases that do not fall under the conditions stated in the Condorcet Jury Theorem for majority voting as well as for groups with low second order reliabilities.

(Hartmann & Rafiee Rad 2018: 1289)

In sum, if Hartmann and Rafiee Rad are right, although majority vote may be more easily implemented as a procedure for solving intragroup disagreement in the case of large groups, deliberating *vis-à-vis* the goal of reaching true collective agreements is roughly as epistemically appropriate as voting by majority rule. This gives an answer to the resolution question. As we will see next, however, this answer is incomplete, since real-life deliberation cases may involve many complexities that make giving a general, straightforward answer to that question a complex matter.

2.2.3 Deliberation in Non-Idealized Conditions

Formal models are surely a great approach to the question of whether deliberation or voting is the most reliable way to solve intragroup disagreement. But deliberation involves many complexities—not present in

voting cases—that have a bearing on its reliability as a collective method for solving intragroup disagreement, and which can make it difficult to give a straightforward answer. To illustrate this, consider Hartmann and Rafiee Rad's model again. As they acknowledge, in order to capture the reliability of more realistic deliberative situations, several assumptions of the model need to be relaxed, such as the assumption that the deliberators are independent, i.e., that the only cause for a group members' verdict is the truth or falsity of the hypothesis in question—whereas other members' verdicts are evidence for the truth or falsity of that hypothesis that don't necessarily break such an independence. Indeed:

- i In real-life deliberation cases, the individual judgments of members of deliberating groups may not be independent from each other.

Or the assumption that the first-order reliability of group members remains unchanged during deliberation. After all:

- ii In real-life deliberation cases, the probability that a given group member is right or wrong about the disputed matter may change along the deliberative process.

Or the assumption that the first-order reliability of group members is independent from their second-order reliability. However:

- iii In real-life deliberation cases, how well a group member estimates how reliable, concerning the disputed matter, other group members are may be influenced by the judgment of those other members.

Other complications that a formal model of deliberation might need to incorporate to better reflect how deliberation compares with voting in real cases include the following. For instance, one crucial assumption of Hartmann and Rafiee Rad's model is that there is full disclosure of the evidence among all group members before deliberation. But:

- iv In real-life deliberation cases, group members may gradually disclose their private evidence to other group members, and may not even disclose any evidence at all.

This is relevant, because in so proceeding deliberating groups can and often times fall prey to *shared information bias*, a tendency to discuss shared evidence, i.e., evidence that most group members possess, in detriment to discussing potentially relevant evidence privately possessed by individual members or only shared by a few of them (cf., Stasser & Titus 1985). If group members have good private evidence, but they do not draw on it during discussion—a situation that is often referred to as a *hidden profile*—the reliability of the deliberative process can be compromised. Several factors can help groups overcome shared information

bias (see §2.3). Interestingly, the kind of groups we are concerned with—those whose goal is to find a correct answer—see this bias diminished by devoting more of their discussions to considering critical clues thus becoming more likely to adopt a correct view when relevant private information remains unshared (Stasser & Stewart 1992).

Another crucial assumption of Hartmann and Rafiee Rad's model is that deliberation proceeds in a series of iterations in which each group member first assigns a probability to the hypothesis in question, then casts a vote, and then each member updates her probability on the basis of the votes of other group members, weighted according to the estimated reliabilities. However:

- v In real-life deliberation cases, group members may discuss the relevant issue one or several times and then take one final single vote to decide which view should stand as the group's view, or simply reach consensus without voting at all.

Other complexities have to do with the different types of evidence distinguished in §2.1. For example, during deliberation group members put their private evidence on the table, which becomes shared evidence. The complication, as we have already pointed out, is that:

- vi In real-life deliberation cases, group members may need to assess two things: how good the evidence shared by other members is (e.g., by judging, among other things, how reliable those members are in gathering good evidence) and how good those other members' assessments of their own shared evidence are (e.g., by judging how reliable they are in assessing the confirmation import of their evidence).

This means that a more realistic model may need to include two measures of second-order reliability, instead of one (see Eder, this volume, for this kind of approach).

In addition, as we have also argued, social evidence—i.e., evidence about the distribution of opinions within the group—can have a defeating effect on its own (i.e., independent of the group's shared evidence) even if it carries no information directly bearing on the question of whether p (like shared evidence does). Relatedly, the very distribution of the disagreement matters and, in particular, when the relevant intra-group disagreement is between a majority and a minority. This is illustrated by extensive research in social psychology on group conformity pressures, and in particular, on *majority influence*. For example, in a famous study by Sherif (1936), subjects were asked to perform a visual task. Subjects whose estimations diverged from those of the majority gradually converged to the latter after being exposed several times to the opinions of the majority. In later studies by Asch (e.g., 1952), the relevant visual task had an obvious correct answer and conformity to the

majority was also observed (although to a lesser extent). Accordingly, it is plausible that:

- vii In real-life deliberation cases, group members who hold a different view to the one held by most members of the group may conform to the majority opinion by repeatedly being exposed to it.

Judging whether a disagreeing majority is right or wrong might be a complex issue. In particular, to judge whether majority influence is epistemically appropriate, one needs to determine whether it is informational (i.e., due to the fact that there is more evidence supporting the relevant opinion) or normative (e.g., due to a desire to fit and avoid social exclusion). Incorporating a corresponding realistic measure of reliability in a formal model of deliberation might accordingly be a complex issue as well.

The issue is even more complex considering the fact that *minorities* also exert influence on majorities. For instance, Moscovici and Zavalloni (1969) observed this kind of effect in a visual task with an obvious correct answer when a minority of subjects gave consistent and unanimous answers that diverged from those of the majority. Thus:

- viii In real-life deliberation cases, group members who unanimously hold a different view to the one held by most members may make the latter conform to their opinions by consistently exposing it to them.

The existence of minority dissent is not necessarily negative at the collective level. Quite the contrary: minority dissent previous to group discussion has been observed to improve the quality of the resulting collective judgments and decisions (e.g., Hightower & Sayeed 1996; Brodbeck et al. 2002; Schulz-Hardt et al. 2005).

Another specific condition widely investigated in social psychology that may affect the reliability of deliberation is the *group polarization* phenomenon (e.g., Stoner 1961; Burnstein & Vinokur 1977; Isenberg 1986)²²:

- ix In real-life deliberation cases, the individual members of like-minded groups may adopt, on average, more extreme views after group discussion than those held before deliberation.

This means, for instance, that in an intragroup disagreement where most group members lean, on average, toward p and only a few toward not- p chances are that if group members discuss whether p or not- p should stand as the group's view, the group's average will more strongly lean toward p . This is a source of collective unreliability, at least in the cases where p is false. Thus, the initial distribution of opinions in a group

featuring an internal disagreement matters for how reliable deliberation is in solving it.

Finally, when it comes to the different kinds of evidence involved in deliberation, the most difficult issue to solve is this:

- x In real-life deliberation cases, it may be unclear what exactly the interplay between the different kinds of evidence (private, shared or social) is, and which one should play a more significant role in whether a group ends up adopting a true or else a false view following deliberation.

As we noted in §2.1, it is an open question which of these three kinds of evidence should have a greater weight in fixing the reliability of individual deliberators. This question might be difficult to address with formalization or empirical research only, and further philosophical investigation is required.

So where does this leave us? Is deliberation as appropriate as a method to solve intragroup disagreement *vis-à-vis* the truth goal as voting is when CJT-style theorems apply? According to Hartmann and Rafiee Rad's model, the answer is 'yes'. However, this answer to the resolution question, albeit on the right track, is not (as they also acknowledge) fully satisfactory: deliberating groups can be affected by a variety of factors that bear negatively (but also positively) on the reliability of deliberation. Some such factors that we have discussed are: (i) the interdependence between the judgments of group members; (ii) changes in their first-order reliability along the deliberative process; (iii) the first-order and second-order reliabilities of group members not being independent; (iv) shared information bias and hidden profile situations; (v) different modes of deliberating, such as several iterations of deliberation and voting, deliberation followed by a single final vote, or deliberation followed by consensus absent voting; (vi) the group members' need to assess the epistemic quality of the evidence shared by others and of the judgments they make about such evidence; (vii) majority influence; (viii) minority influence; (ix) group polarization; and (x) the complex interplay between private, shared, and social evidence. If anything involves complexity, that is the question of whether deliberation is a truth-conducive method for solving intragroup disagreement.

2.3 Assessing for Evidence

Truth is not the end of the story, however. It is not unusual that members of a group featuring an internal disagreement are not merely interested in settling on a true collective view, but on a view that is supported by the best evidence individually possessed by them. Of course, the truth and evidence goals are not incompatible and are in fact oftentimes