

The top half of the book cover features a teal background with a close-up, slightly out-of-focus image of a pair of glasses. The frames and lenses are visible, creating a pattern of overlapping circles and lines. In the top left corner, there is a solid teal square containing the word "OXFORD" in white, uppercase letters.

OXFORD

Against Coherence

Truth, Probability, and Justification

ERIK J. OLSSON

Against Coherence

This page intentionally left blank

Against Coherence

Truth, Probability, and Justification

ERIK J. OLSSON

CLARENDON PRESS • OXFORD

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford ox2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

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

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

© Erik J. Olsson 2005

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

First published 2005

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

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

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloging-in-Publication Data

Olsson, Erik J.

Against coherence : truth, probability, and justification / Erik J. Olsson.
p. cm.

Includes bibliographical references and index.

ISBN 0-19-927999-3 (alk. paper)

1. Truth—Coherence theory. 2. Skepticism. I. Title.

BD171.O47 2005 121—dc22 2004029060

Typeset by Newgen Imaging Systems (P) Ltd., Chennai, India
Printed in Great Britain
on acid-free paper by
Biddles Ltd., King's Lynn, Norfolk

ISBN 0-19-927999-3 978-0-19-927999-9

1 3 5 7 9 10 8 6 4 2

*Dedicated to
Maryam and Nina,
wife and daughter,
with love and appreciation*

This page intentionally left blank

Preface

Coherence has to do with the degree to which items ‘hang together’, ‘dovetail’, or ‘mutually support’ each other. This book is about coherence and truth, its aim being to determine whether there is any substantial connection between these two concepts. Is a system that is coherent thereby highly likely to be true? Is a system that is *more* coherent than another system thereby *more* likely to be true? These questions will be central to our endeavours.

Why should we care about coherence and truth in the first place? One reason has to do with scepticism. If we can show that our ordinary beliefs are highly likely to be true, then presumably we are justified in holding on to them. A persistent belief in the history of philosophy, one that still has distinguished adherents, is that we should be able to conclude something of this nature by inspecting the degree in which our beliefs cohere. Many take this position, I speculate, because it is perceived to be the anti-sceptic’s last resort. In ascertaining the extent to which coherence can justify our beliefs ‘from scratch’ this book is partly a contribution to the philosophical debate over radical scepticism.

At the same time, and no less importantly, this essay is intended to be a general contribution to the probabilistic study of common sense and scientific reasoning. When we hear the same story reported twice by different sources, we are inclined to believe what is being said, even if we initially did not attach a very high credibility to each reporter taken singly. Such coherence reasoning is, in the words of an early theorist, ‘immanent in all our thinking’ (Ewing 1934: 231). This essay attempts to systematize such thought processes by bringing them under one probabilistic hat, thus making possible a precise investigation of the relation between coherence and (likelihood of) truth in normal, non-sceptical contexts. There is a non-negligible unification bonus associated with this project, as it connects the study of coherence with probabilistic work in philosophy of law, philosophy of religion, and confirmation theory, not to mention the

intimate connection that emerges with artificial intelligence and its Bayesian networks.

As is shown in this book, coherence can have a dramatic impact on the likelihood of truth. Agreeing items of information that are relatively improbable, when taken singly, can be practically certain, when combined. The main insight which I hope can be derived from this book is that the connection between coherence and truth is, nonetheless, too weak to allow coherence to play the role it is supposed to play in a convincing response to radical scepticism. This is so for purely probabilistic reasons that, I hope, everyone could at least in principle agree upon, quite independently of more controversial matters like the externalism–internalism controversy or the philosophical interpretation of probability statements. A further puzzle on which this essay attempts to shed light has its roots in the fact that coherence theorists have been unable to reach anything like a consensus on how to define their central notion. This book explains why this is so: coherence is in a sense not definable. More precisely, there is no way to specify an informative notion of coherence that would allow us to draw even the minimal conclusion that more coherence means a higher likelihood of truth other things being equal (in favourable circumstances). This strongly suggests that the notion of a ‘degree of coherence’ is itself an incoherent one. This does not mean, of course, that we should not, if the circumstances are favourable, assign a relatively high probability to that which is agreed by the many. It does mean, however, that there seems to be little to say about coherence and truth in positive terms beyond this epistemologically inconsequential item of common sense. It is still possible for incoherence to play an important negative, falsificatory role in reasoning. This more constructive line of thought is explored in the final chapter.

This book grew out of research I did on coherence and probability at the University of Constance, Germany, during the period 1999–2003, and I am grateful to a number of people for supporting this work in various ways, especially to Hans Rott, Wolfgang Spohn, André Fuhrmann, and Ulf Friedrichsdorf. Ludwig Fahrbach’s comments on the different versions of my ideas were also valuable. Christopher von Bülow proof-read a previous version of the whole manuscript. My wife Maryam, who holds a Ph.D. in Physics, helped me with some of the mathematics in Chapter 8.

I recall useful communications with many fine researchers, especially Keith Lehrer, Isaac Levi, Paul Thagard, and Carl G. Wagner. I am indebted as well to my new colleagues in Lund, Sweden, in particular to Wlodek Rabinowicz and Johan Brännmark, for their input.

In the process of writing this book I have benefited a lot from my previous joint work with Luc Bovens and with Tomoji Shogenji, two able analytic philosophers in the younger generation who happen to share my fascination with the problems of coherence and probability. As will be clear to the reader of Parts II and III of this book, though, I strongly disagree with some of the conclusions they have reached in their other work.

Hans Rott, Wolfgang Spohn, and Luc Bovens wrote detailed referee reports on an earlier version of this book which was submitted as a second (habilitation) thesis at the University of Constance. I am much obliged to them for giving me access to their acute criticisms and many constructive proposals, which led to several substantial revisions of the original manuscript. I am also greatly indebted to two anonymous Readers at Oxford UP for their criticism, which forced me to think harder about several issues.

Bits and pieces of what was to become this book have been presented to North American audiences at the universities of Arizona, Columbia (New York), Columbia (Missouri), Miami, Pittsburgh, and Waterloo (Canada), and to German audiences in Constance, FU Berlin, Bielefeld, and Leipzig. I am grateful to these audiences for their input and criticism. Finally, I would like to take the opportunity to thank the German Research Council (Deutsche Forschungsgemeinschaft) for financing my research during my period in Constance and to the Swedish Research Council (Vetenskapsrådet) for funding my present research, including the completion of this book.

E. J. O.

Acknowledgements

While some of the results and arguments presented in this book have been published elsewhere, most of the material has not. Essentially new are Chapters 2, 4, 5, 8, and 9. New also is the impossibility theorem that is stated and discussed at the end of Chapter 7 and proved in Appendix B. Most other chapters contain lots of new material in addition to the published works they are based on. I should add that many of the arguments already published have undergone extensive reconstruction and, I hope, improvement. In Chapter 3, the discussion of C. I. Lewis's problem of fixing the individual credibility is an extension of a line of reasoning that appeared at the end of my 'What Is the Problem of Coherence and Truth?', *Journal of Philosophy*, 99 (2002), 246–72. This reasoning was refined in a joint paper with Tomoji Shogenji entitled 'Can We Trust our Memories? C. I. Lewis's Coherence Argument' (forthcoming in *Synthese*) of which I have also taken advantage. Chapter 6 and Chapter 7 are based on the first part of the *Journal of Philosophy* paper and on 'Why Coherence is not Truth-Conducive', *Analysis*, 61 (2001), 236–41. They also draw on my involvement in a debate with Luc Bovens and his associates Brandon Fitelson, Stephan Hartmann, and Josh Snyder about concurring testimonies. My contributions to that debate appeared as 'Corroborating Testimony, Probability and Surprise', *British Journal for the Philosophy of Science*, 53 (2002), 273–88; and, in the same issue (565–72), 'Corroborating Testimony and Ignorance: A Reply to Bovens, Fitelson, Hartmann and Snyder'. The discussion of the Klein–Warfield argument in Chapter 6 and the proof in Appendix A draw upon a joint paper with Luc Bovens which appeared as 'Believing More, Risking Less: On Coherence, Truth and Non-trivial Extensions', *Erkenntnis*, 57 (2002), 157–250. The short section on Lehrer in Chapter 9 is based on my introduction to *The Epistemology of Keith Lehrer*, a book I edited that was published by Kluwer in 2003. Chapter 10 is based on the article 'Lassen wir den Skeptiker nicht zu Wort kommen: Pragmatismus und radikaler Zweifel', in *Pragmatisch denken*, a volume edited by André Fuhrmann and myself which was published by Ontos Verlag in 2004.

Contents

1. Introduction	I
Part I Does Coherence Imply Truth?	
2. Coherence, Truth, and Testimony	9
2.1 Why Coherence?	9
2.2 Coherence—an Elusive Concept	12
2.3 Pinning down the Coherence Theorist	16
2.4 Truth and Agreement	21
2.5 A Simple Witness Model	24
2.6 Conditions for Convergence	26
2.6.1 Testimonial Independence	27
2.6.2 Individual Credibility	30
2.7 Convergence Parameters	31
2.8 Challenges for the Coherence Theorist	32
3. C. I. Lewis's Radical Justification of Memory	34
3.1 The Problem of Justifying Memory	34
3.2 Lewisian Witness Scenarios	35
3.2.1 Lewis on Witness Corroboration	36
3.2.2 Lewis Scenarios and Surprising Agreement	38
3.2.3 Modelling a Lewis Scenario	39
3.3 Lewis on the Convergence Conditions	48
3.3.1 The Transcendental Argument for Individual Credibility	49
3.3.2 The Verificationist Argument for Individual Credibility	52
3.4 The Individual Credibility 'need not be Assigned'	55
3.5 A Note on Lewis's Definition of Independence	58
4. Laurence Bonjour's Radical Justification of Belief	61
4.1 The Problem of Justifying Beliefs	61
4.2 Bonjour on Justification from Scratch	66
4.3 Lying and Individual Credibility	69
4.4 Coordinated Lying and Independence	72
4.5 Consequences for Bonjour's Anti-scepticism	74

5.	C. A. J. Coady's Radical Justification of Natural Testimony	77
5.1	The Problem of Justifying Natural Testimony	77
5.2	Hume's 'Fatal Ambiguity'	79
5.3	The Argument for Individual Credibility	81
5.4	The Invocation of Cohesion and Coherence	84
5.5	Coady's Fatal Ambiguity	88
5.6	Closing Remarks on the Anti-sceptical Use of Coherence	90

Part II Does More Coherence Imply Higher
Likelihood of Truth?

6.	Making the Question Precise	95
6.1	Degrees of Coherence	95
6.2	Coherence and Logical Closure	101
6.3	Testimonial Truth Conduciveness	102
6.4	Why the Klein–Warfield Argument Fails	105
6.5	Testimonial vs. Explanatory Coherence	108

7.	A Negative Answer	112
7.1	Independence and Individual Credibility	112
7.2	The Need for a <i>Ceteris Paribus</i> Clause	116
7.3	Should Specificity (Strength) be Held Fixed?	119
7.4	Should Size be Held Fixed?	123
7.5	L. J. Cohen on the Influence of the Prior on the Posterior	125
7.6	The Impossibility of Coherence	134

Part III Other Views

8.	How not to Regain the Truth Connection: A Reply to Bovens and Hartmann	143
9.	Other Coherence Theories	156
9.1	Nicholas Rescher	156
9.2	Donald Davidson	157
9.3	Keith Lehrer	159
9.4	Paul Thagard	162

Part IV Scepticism and Incoherence

10.	Pragmatism, Doubt, and the Role of Incoherence	173
10.1	Cartesian Scepticism	173
10.2	Jamesian Wagering	176

10.3	Peirce's Reply to Scepticism	187
10.4	More on Incoherence as a Reason for Doubt	190
10.5	Three Roads to Scepticism	193
10.6	Comparison with Other Contemporary Responses	198
10.7	Conclusion	203
Appendix A: Counter-example to the Doxastic Extension Principle		205
Appendix B: Proof of the Impossibility Theorem		211
Appendix C: Proofs of Observations		216
<i>References</i>		224
<i>Index</i>		229

This page intentionally left blank

Introduction

Does coherence imply truth? This is our central problem, and one could in principle imagine many possible ways of attacking it. While it may be implausible to think that a system that is coherent is thereby guaranteed to contain only true propositions, it is conceivable that coherence could imply verisimilitude, so that a system, in virtue of being coherent, is at least close to the truth. The problem of coherence and truth, like so many other philosophical issues, cannot be exhaustively examined in one single book of manageable length. Every author must accordingly make a selection of aspects he or she wants to study within the boundaries of one volume. This book is about the possibility of a *probabilistic* connection. Is a system that is coherent thereby highly likely to be true? Is a system that is *more* coherent than another system thereby *more* likely to be true?

There is a further decision to be made, one that seems no less unavoidable. One could imagine many ways to explicate the central notion of coherence, and in the literature one finds accordingly a number of different suggestions—ranging from purely ‘logical’ definitions in terms of logical consistency or mutual derivability to definitions that make use of neural network concepts. Some of these proposals are implausible, others perhaps not. In making the first decision we also, in effect, made the second. Given the choice to focus our attention on possible probabilistic relations between coherence and truth, the natural further decision is to zoom in on concepts of coherence that are explicable in similar probabilistic terms. I should add that I do discuss other conceptions in Chapter 9, though relatively briefly.

The works that have proved to be most relevant to the aims of this essay are those belonging to what might be called the ‘Harvard school’ of coherence theorists, the most important single contributions to

this tradition of thought being C. I. Lewis's *A Theory of Knowledge and Valuation* from 1946 and Laurence BonJour's *The Structure of Empirical Knowledge* from 1985. Lately, this sort of theory and its supposed 'truth conduciveness' claim have been the subject of intense debate in the journal *Analysis*, a controversy that was ignited by a thought-provoking paper by Peter Klein and Ted A. Warfield in which an attempt was made to show, by counter-example, that coherence is not, in fact, truth conducive. Also highly relevant to my concerns have been L. Jonathan Cohen's *The Probable and the Provable* (1977) and C. A. J. Coady's *Testimony: A Philosophical Study* (1992). Parts I–III of the present book are the results of my efforts to provide systematic answers to the problems addressed in this core literature.

The first part of this book examines the thesis that coherence implies truth. Having criticized the various definitions of coherence in the literature for being either too vague or plainly inadequate, I argue, in an attempt to pin down the coherence theorist, that full agreement among testimonies must be regarded as a case of coherence. This is important since it opens up the possibility of putting the coherence theorist's canon to the test. Can we at least show that full agreement implies a high likelihood of truth? As I go on to argue, with reference to a simple witness example, the answer to that question seems to be in the negative. First of all, the standard situations in which the addition of an agreeing testimony has a positive effect on the likelihood of truth are such that the reports satisfy the further conditions of being collectively independent and individually to some degree credible. And, what is more, even under such favourable circumstances, the effect of adding one more agreeing testimony on the likelihood of truth need not be very impressive, since the latter depends on the prior probability of what is being reported and also on the exact degree of credibility each witness has taken singly. So, far from guaranteeing a high likelihood of truth *by itself*, testimonial agreement can apparently do so only if the circumstances are favourable as regards independence, prior probability, and individual credibility. These are troublesome facts for Lewis and BonJour, who want to justify beliefs or memories from scratch using coherence reasoning and who have to show, therefore, that these facts can somehow be accommodated by their theories. I argue that neither Lewis nor BonJour is capable of meeting these challenges convincingly.

The main problem for Lewis is his insistence that coherence implies a likelihood of truth sufficient for ‘rational and practical reliance’, given independence and some positive individual credibility, regardless of the specific degree of individual credibility pertaining to the reports. One weak point in Bonjour’s reasoning is his reliance on the thesis that coherence can guarantee a high likelihood of truth even if the independent reports have no individual credibility at all. The upshot of all this is that the coherence theorist has the choice of either obscuring the concept of epistemological coherence, by dissociating it from testimonial agreement, or rejecting the idea that coherence implies a high likelihood of truth, where the latter alternative stands out as the far more reasonable choice. At the end of Part I, I argue that Coady’s recent coherence-based attempt to provide a radical justification of our trust in the word of others (‘natural testimony’) is structurally identical to Lewis’s proposed vindication of memory and that therefore my criticism of Lewis applies to Coady’s theory as well. I also point out certain other difficulties in Coady’s reasoning.

Part II is a systematic attempt to come to grips with the central problem in the aforementioned *Analysis* debate, and its focus is more on normal uses of coherence than on anti-sceptical uses. If we cannot say that coherence implies a high likelihood of truth, can we at least say that coherence is truth conducive in the sense that *more* of it implies a *higher* likelihood of truth? Whereas it was sufficient for the purposes of Part I to consider cases of full agreement, the comparative question obviously presupposes a comparative conception of coherence that allows us to determine when there is more or less of it. Along the way, I take the opportunity to say why, exactly, Klein and Warfield’s controversial counter-example to the truth conduciveness of coherence fails. Having clarified the question, I move on to search for some answers. It is pointed out that coherence cannot be truth conducive in the comparative sense in the absence of the conditions of independence and individual credibility. Furthermore, there is a need for a *ceteris paribus* clause: more coherence can imply a higher likelihood of truth only if all other things are equal. This leads me to the remaining question: are there any (interesting) measures of coherence that are truth conducive *ceteris paribus* given independence and individual credibility? This question is

closely connected to an issue raised by L. Jonathan Cohen about the role of specificity in the probabilistic theory of concurring testimony. The discussion of Cohen is interesting in its own right but for our purposes its main import lies in the fact that it leads up to an impossibility result: there cannot be a non-trivial coherence measure that is truth conducive *ceteris paribus* in the sort of witness scenarios that coherence theorists have typically taken interest in. This puts us in an excellent position to explain why coherence theorists have been unsuccessful in defining their central notion: coherence is in a sense not definable.

Part III is devoted to some remaining issues. As Luc Bovens and his colleagues have pointed out, in a recent debate with me, Cohen's claims about specificity can, in some degree at least, be saved if one is prepared to invoke the Principle of Indifference. In this part of the book, this general strategy is examined with respect to the prospects of deploying it for the following purposes: (a) modelling realistic witness scenarios, (b) solving Lewis's problem regarding individual credibility, and (c) avoiding the aforementioned impossibility result. I argue that the strategy fails on all three accounts. Another issue addressed here is how my account of coherence and truth compares with other well-known approaches. I discuss briefly in this connection the coherence theories of Nicholas Rescher, Donald Davidson, Keith Lehrer, and, in some more detail, Paul Thagard.

In Part IV, finally, I return to the issues of scepticism and radical doubt. Rebutting scepticism is not the central aim of this book which is, first and foremost, a critique of coherence theories. Still, given that coherence is not the answer to scepticism, one may wonder what a more adequate response would look like. The starting point of the whole discussion is the tacit assumption in much literature on scepticism that we must accept the sceptic's challenge or face charges of irrationality. Thus, the fear arises that philosophical and epistemological breakdown is imminent unless one somehow manages to justify one's beliefs, memories, and other commitments on neutral ground. This perceived pressure may, I speculate, be part of the explanation why philosophers are led to advocate anti-sceptical theories that upon serious examination leave much to be desired, coherence theories being only one case in point. In this part, I examine the reasons for engaging in radical justification in the first place. Already in Part I,

I discuss briefly the moderate pragmatism of C. I. Lewis and his ultimate appeal to pragmatic considerations in his justification of the individual credibility of our memories. The upshot of the argument launched here is that, in the end, a version of pragmatism can provide us with a compelling reply to the radical sceptic. My point of departure is William James, whose response is seen to consist in the application of a wager argument to the sceptical issue in analogy with Pascal's wager. The strategy of C. S. Peirce, on the other hand, amounts to a direct rejection of one of the sceptic's main premisses: that we do not know we are not deceived. I argue that while the Jamesian attempt is ultimately untenable, Peirce's argument contains the core of a convincing pragmatic rebuttal of scepticism.

The Peircean discussion reveals the second reason why returning to the topic of scepticism is appropriate in this context, for it gives a new perspective on coherence—or, rather, *incoherence*. Following Peirce, I argue that genuine doubt is always preceded by some sort of incoherence. The suggestion then is that whereas the lack of a substantial connection between coherence and truth makes it severely problematic to claim that coherence has a role to play in the process whereby beliefs are acquired or justified, it can still be maintained that incoherence is the driving force in the process whereby beliefs are retracted. On this proposal, which is here tentatively explored, the role of coherence in our enquiries is negative rather than positive.

This page intentionally left blank

Part I

Does Coherence Imply Truth?

This page intentionally left blank

Coherence, Truth, and Testimony

2.1 Why Coherence?

In ordinary life we usually rely on the information sources that we have at our disposal, placing our trust in the testimony of other people as well as in the testimony of the senses. Such reliance, as a number of authors have pointed out, is automatic and routine.¹ This is most obvious for the testimony of the senses. Thus, I come to believe that my friend is over there as the direct effect of observing him without in any way inferring his presence from other beliefs I have. But the same is basically true of testimonies from other people. If the secretary tells me that my colleague was in his office just a moment ago, I simply believe it.

While the reception of testimony from various sources is normally unreflective, it is not thereby uncritical. Testimony is accepted so long as there is no explicit reason to doubt the credibility of the reporter, i.e. so long as certain trouble indicators are not present. The mechanism is deactivated if, for instance, we find positive reasons to question the motives of our informant. Is she trying to deceive us? Even an informant with the best of intentions may turn out not to be trustworthy if there are signs that she acquired her information under problematic circumstances (e.g. under bad lighting conditions). If there are no special reasons for caution, the unreflective mechanism of reliance is invoked and one single testimony suffices to settle the matter, at least for the time being.²

¹ See, for example, Levi (1991) and Coady (1992: 47). See also my recent debate with Levi about how to construe reliance more precisely in Olsson (2003a) and Levi (2003).

² Cf. Coady (1992: 47): ‘We may have “no reason to doubt” another’s communication even when there is no question of our being gullible; we may simply recognize that the standard warning

Coherence becomes relevant once the reliability of our informants is, for some reason, in doubt, so that we are unable to take that which is being reported at face value. In this case it may pay off to listen to more than one source. If the sources cohere or agree to a large extent in their reporting we may conclude that what they say is true, even though this conclusion could not have been reached as the effect of listening to one of the sources only. If, for instance, the first dubious witness to be queried says that John was at the crime scene, the second that John has a gun, and the third that John shortly after the robbery transferred a large sum to his bank account, then the striking coherence of the different testimonies would normally make us pretty confident, notwithstanding their individual dubiousness, that John is to be held responsible for the act.

C. I. Lewis made the same point when he asked us to consider a case of ‘relatively unreliable witnesses who independently tell the same circumstantial story’ (1946: 246).³

For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon.

The resulting probability of what is agreed need not merely be high but may even suffice for practical certainty:

Take the case of the unreliable observers who agree in what they report. In spite of the antecedent improbability of any item of such report, when taken separately, it may become practically certain, in a favorable case, merely through congruent relations to other such items, which would be similarly improbable when separately considered. (352)

As Lewis makes clear, the foregoing remarks apply not only to witness reports but quite generally to ‘evidence having the character of “reports” of one kind or other—reports of the senses, reports of memory, reports of other persons’ (347). Take, for instance, memory reports:

[S]omething I seem to remember as happening to me at the age of five may be of small credibility; but if a sufficient number of such seeming

signs of deceit, confusion, and mistake are not present. This recognition incorporates our knowledge of the witness’s competence, of the circumstances surrounding his utterance, of his honesty, of the consistency of the parts of his testimony, and its relation to what others have said, or not said, on the matter.’

³ Throughout this book, all references to Lewis concern his 1946 essay *Knowledge and Valuation*.

recollections hang together sufficiently well and are not incongruent with any other evidence, then it may become highly probable that what I recollect is fact. It becomes thus probable just in measure as this congruence would be unlikely on any other supposition which is plausible. (352)

Throughout this book I will take ‘testimony’ in the wide sense to include not only witness testimony but also, for instance, the ‘testimony of the senses’ and the ‘testimony of memory’. Thus, I use ‘testimony’ in the same sense in which Lewis uses ‘report’. I will sometimes employ Coady’s term ‘natural testimony’ to refer to actual assertions (by witnesses etc.).⁴

The foregoing remarks are intended to highlight the *normal* use of coherence, i.e. its employment in enquiries characterized by (1) some of the warning signs being present making it impossible to accept testimonies at face value, but (2) there being nonetheless a substantial body of background assumptions upon which we can, in fact, rely. Our background information, which is not in doubt in the context of the given enquiry, may tell us, for instance, that the informants are independent of each other and that they, while falling short of full reliability, are nonetheless to be regarded as relatively reliable.

What is especially striking about coherence reasoning is that by combining items of information which are in themselves almost worthless one can arrive at a high probability of what is being reported. Indeed, it is salient how little knowledge of the reporters seems necessary for coherence to result in high likelihood of truth. We can, it seems, be *almost* entirely ignorant about the quality of our reporters and still arrive at practical certainty as the effect of observing their agreement. At least, this is what Lewis seems to suggest.

There is but a small step from arguing that coherence works under almost total ignorance to holding that it does so even if we remove the ‘almost’. If coherence is so successful in coping with context where *very little* is taken for granted, could it not also be invoked where *nothing* is? Hence the *anti-sceptical* use of coherence, i.e. the employment of coherence reasoning in sceptical contexts. These contexts are characterized by everything being called into question,

⁴ For an account of different senses of ‘testimony’, including its use in legal contexts, see chapter 2 in Coady (1992).

except facts of a mere report character. The allowed reports typically state that a person believes or remembers this or that. The claim, then, is that a person can, using coherence reasoning, legitimately recover her trust in her beliefs or memories from this meagre base. We can, it is contended, start off with literally nothing—as the sceptic insists—and yet, upon observing the coherence of our *de facto* memories or beliefs, conclude that those memories or beliefs are highly likely to be true.

Thus we are led to the kind of coherence theory advocated by C. I. Lewis and Laurence Bonjour. Their theories differ on interesting points, as we will see later, but the general concept is the same: both intend to provide a final validation—or, as I will also say, a *radical justification*—of our empirical knowledge through the anti-sceptical use of coherence reasoning on initially highly dubious data in the form of mere reports on what we believe or (seem to) remember. Their anti-sceptical theories are partly based on certain claims about what is supposed to be true of witness cases, typically accentuating the supposed success of coherence reasoning in such cases. These claims are then said to apply equally to various sceptical scenarios involving beliefs or memories.

Given our interest in the relation, if there is any, between coherence and truth, there are two questions in need of detailed answers. First, what connection, if any, is there between these concepts in normal contexts and, second, what is there to say about the supposed relation in sceptical contexts? In addressing the first issue, this book is a contribution to the probabilistic study of normal common sense and scientific employment of coherence reasoning. In attending to the second concern, it also has implications for the philosophical debate over radical scepticism.

2.2 Coherence—an Elusive Concept

Before we can get anywhere with the question whether coherence implies truth we must overcome a serious obstacle. As Davidson observes, truth is ‘beautifully transparent’ in comparison to coherence (1986: 309), and so we must find some way of assigning definite meaning to the latter. Unfortunately, the literature is surprisingly

silent in this regard, apart from some very general and vague characterizations of coherent sets in terms of ‘mutual support’, of their elements ‘hanging together’, and so on. Not surprisingly, it has become a standard objection to coherence theories that their advocates fail to provide a detailed account of the central notion, thus reducing it ‘to the mere uttering of a word, coherence, which can be interpreted so as to cover all arguments, but only by making its meaning so wide as to rob it of almost all significance’ (Ewing 1934: 246). As Nicholas Rescher, another prominent coherence theorist, puts it, ‘the coherence theorists themselves have not always been too successful in explicating the nature of coherence’ (1973: 33).

The absence of a clear account has been noted as a troublesome fact ever since the days of the British idealists, and more recent coherence theories fare no better, in the lights of their critics, than their idealist ancestors did. Thus, Marshall Swain, referring to Laurence Bonjour’s celebrated coherence theory as put forward in his 1985 book, complains that ‘[o]ne of the most disappointing features of Bonjour’s book is the lack of detail provided in connection with the central notion of coherence’ (1989: 116).

In the few cases where coherence theorists have actually proposed clear definitions, they can be seen, on closer scrutiny, to be incorrect. One case in point is A. C. Ewing’s restrictive definition of coherence in terms of mutual logical entailment. A set of propositions is said to be coherent if its elements are ‘so related that any one proposition in the set follows with logical necessity if all the other propositions in the set are true’ (1934: 229).⁵ But there is an obvious way to relax it by allowing weaker relations than logical entailment to be coherence inducing. C. I. Lewis, accordingly, defines coherence—or, to use his favoured term, ‘congruence’—as follows (338):⁶

A set of statements, or a set of supposed facts asserted, will be said to be congruent if and only if they are so related that the antecedent probability of any one of them will be increased if the remainder of the set can be assumed as given premises.

⁵ Ewing’s final proposal is slightly more complicated in ways that do not affect the point I am making here.

⁶ Lewis prefers ‘congruence’ to ‘coherence’ because he wants to mark his departure from British post-Kantian idealism and its ‘coherence theory of truth’ (1946: 338).

The set S consisting of A_1, \dots, A_n is congruent relative to a probability distribution P just in case $P(A_i/B_i) > P(A_i)$ for $i = 1, \dots, n$, where B_i is a conjunction of all elements of S except A_i .

But as the following example shows, Lewis's definition is also flawed.⁷ Suppose there to be a reasonable number of students and a reasonable number of octogenarians (80–89-year-olds). Suppose that all and only students like to party, that all and only octogenarians are birdwatchers, and that there are some, but very few, octogenarian students. A murder happened in town. Consider the following propositions:

- A_1 = 'The suspect is a student'
- A_2 = 'The suspect likes to party'
- A_3 = 'The suspect is an octogenarian'
- A_4 = 'The suspect likes to watch birds'

The set $\{A_1, A_2, A_3, A_4\}$ is congruent in Lewis's sense: each individual proposition is made more probable by assuming the others to be true. And yet the set is intuitively anything but coherent, one half of the story (the one about the partying student) being highly unlikely given the other half (the one about the birdwatching octogenarian), and vice versa. We note that Lewis's definition, though incorrect in general, can still be adequate in the two-proposition case.

Beside Lewis's often quoted definition, the following 'coherence criteria' due to BonJour are generally taken to reflect the state of the art (1985: 95–9):

1. A system of beliefs is coherent only if it is logically consistent.
2. A system of beliefs is coherent in proportion to its degree of probabilistic consistency.
3. The coherence of a system of beliefs is increased by the presence of inferential connections between its component beliefs and increased in proportion to the number and strength of such connections.
4. The coherence of a system of beliefs is diminished to the extent to which it is divided into subsystems of beliefs which are relatively unconnected to each other by inferential relations.

⁷ This counter-example is adopted from Bovens and Olsson (2000: 688–9).

5. The coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system.

While this account has many merits, it is also unclear on several crucial points. For one, how are we to measure the number and strength of inferential connections? How can we assess whether sub-systems are ‘relatively unconnected’ or not? What is the connection between coherence as an absolute notion (see the first criterion) and coherence as a matter of degree (see the other criteria)? It is difficult to see how we could get anywhere with our investigation into coherence and truth unless at least some of these questions are given clear answers. Another problem, which BonJour’s theory shares with multi-aspect theories generally, is whether the different aspects are really independent of each other, or whether in fact some coherence criteria are rendered obsolete in the presence of the others.

Most seriously, it is even far from obvious that the criteria are correct. Take the second one for instance. What reasons do we have for thinking that a system of beliefs should be coherent in proportion to its degree of probabilistic consistency? The notion of ‘probabilistic inconsistency’ is elucidated as follows:

Suppose that my system of beliefs contains both the belief that P and also the belief that it is extremely improbable that P . Clearly such a system of beliefs may perfectly well be logically consistent. But it is equally clear from an intuitive standpoint that a system which contains two such beliefs is significantly less coherent than it would be without them and thus that probabilistic consistency is a second factor determining coherence. (ibid.: 95)

But this cannot be true in general. Having bought a ticket for the National Lottery, I watch the TV and learn that my ticket is among the winners. It is my number that is being displayed there on the screen; there is no question about it. And yet I also know that before the event occurred it was extremely improbable that my ticket should win. If BonJour’s second criterion were correct, my believing that my ticket has won should reduce the coherence of my belief system. But *pace* BonJour this is clearly counter-intuitive. It is absurd to think that lottery winners are generally slightly incoherent in believing that they have won—even if their belief relies on absolutely reliable

evidence. I will return to the role of incoherence and anomaly in Chapter 10.⁸

2.3 Pinning down the Coherence Theorist

These problems notwithstanding, is there anything that coherence theorists should be able to agree on as to the nature of coherence, apart from the vague idea of coherence being determined by connections between beliefs? Take a case of two witnesses, Smith and Jones, testifying individually to the effect that another man, Forbes, has committed a certain crime. Now if this is not a case of coherence, then, I must confess, I have no idea of what that notion could possibly involve. After all, the witnesses say exactly the same thing, and so what they say could hardly be in greater ‘harmony’, exhibit greater ‘mutual support’, or ‘hang better together’, to refer to some of the usual characterizations of coherent sets. Not allowing cases of testimonial agreement to be cases of coherence is indeed committing the very fallacy that Ewing warned us of, that of ‘robbing coherence of all significance’.

It is useful to distinguish three progressively stronger claims about coherence and testimonial agreement:

1. Coherence as well as incoherence can be applied meaningfully to cases of testimonial agreement without any category mistake thereby being committed.
2. Cases of testimonial agreement are also cases of coherence.
3. Testimonial agreement is more than just coherent; it is very coherent.

Obviously, (3) implies (2). The latter, moreover, entails (1): if agreement among testimonies is properly described as a coherent situation,

⁸ Bonjour adds the following: ‘First, it is extremely doubtful that probabilistic inconsistency can be entirely avoided. Improbable things do, after all, sometimes happen, and sometimes one can avoid admitting them only by creating an even greater probabilistic inconsistency at another point. Second, probabilistic consistency, unlike logical consistency, is plainly a matter of degree, depending on (a) just how many such conflicts the system contains and (b) the degree of improbability involved in each case (ibid.).’ In Chapter 10 I will argue that believing that some improbable thing has happened induces incoherence only if there is an alternative explanation of the event, an explanation that would make it very likely to occur.

it cannot be a category mistake to apply the concept of coherence in such cases. I will proceed to argue that all three claims can be plausibly attributed to both Lewis and Bonjour.

Let us start with Lewis. When illustrating the capacity of coherence to raise probability, Lewis suggests, as we saw, that we consider relatively unreliable witnesses who independently tell the same circumstantial story, in which case ‘congruence of the reports establishes a high probability of what they agree upon’ (346). Evidently, telling the same story is for Lewis a case of coherence or, as he prefers, congruence. So, Lewis’s remarks on witness cases commit him to (2) and, by implication, also to (1).

But does (2) really follow from Lewis’s definition of congruence? This is a surprisingly subtle matter. The difficulty concerns how exactly to conceive of the sets to which the concept of coherence is supposed to be applicable. Lewis suggests that the elements of those sets are ‘supposed facts asserted’. Which are the supposed facts asserted in the Forbes case? They are ‘Forbes did it’ as asserted by Smith and ‘Forbes did it’ as asserted by Jones. Hence, the set of supposed facts asserted is {‘Forbes did it’, ‘Forbes did it’}. But this set is identical with the singleton {‘Forbes did it’}. Now combine this observation with what I will refer to as *Rescher’s Principle*, according to which ‘[c]oherence is . . . a feature that propositions cannot have in isolation but only in groups containing several—i.e. at least two—propositions’ (Rescher 1973: 32) and it follows that the set of supposed facts asserted in the Forbes scenario is neither coherent nor incoherent. To say otherwise would be to commit a category mistake, and so (1) is violated. This is in conflict with Lewis’s commitment to (1) in his examination of witness cases.

Suppose instead that we choose to interpret Lewis as intending to apply congruence not to sets of supposed facts asserted but to sets of assertions of supposed facts.⁹ There are, in the Forbes case, two assertions of supposed facts: Smith’s assertion that Forbes did it and Jones’s assertion to the same effect. The set of assertions of supposed facts is accordingly {‘Smith says that Forbes did it’, ‘Jones says that Forbes did it’}. This set is not a singleton, and so Rescher’s Principle does

⁹ This is one way of interpreting Shogenji’s recommendation that, for the purposes of coherence evaluation, beliefs should be individuated by their sources rather than by their contents (2001: 150).

not apply and we have no reason to believe that (1) is violated. But does (2) hold? Is this set of assertions a coherent set? Let us use Lewis's own definition (which, I submitted, is relatively unproblematic in the two-proposition case) to settle the matter. Does the one assertion in the set raise the expectation that the other is true? There is no simple yes-or-no answer to that question. What the answer is depends on several factors about which our example remains silent. Plausibly, Smith's testifying against Forbes would raise our expectations that Jones would too, provided Smith and Jones are individually somewhat reliable and collectively independent. It would not if, for example, Smith is a highly reliable expert witness and Jones a notorious liar, in which case Smith's testifying against Forbes would in fact reduce the probability of Jones's testifying to the same effect. For, if Smith, the expert, says that Forbes did it, then probably he did. Hence, Jones, the liar, will probably testify that he did not. On this alternative construal of Lewis's 'supposed facts asserted', (2) will not be true in general, and what we have is, again, a conflict between Lewis's definition of congruence and his specific discussion of agreement at other places.¹⁰

We can make sense of Lewis by interpreting his sets not as unordered but as *ordered* sets. Hence a set of supposed facts asserted is not an entity of the type $\{A_1, A_2, \dots, A_n\}$ but one of the type $\langle A_1, A_2, \dots, A_n \rangle$. In the Forbes case, for instance, the set of supposed facts asserted is not $\{\text{'Forbes did it'}, \text{'Forbes did it'}\} = \{\text{'Forbes did it'}\}$, which is a singleton, but $\langle \text{'Forbes did it'}, \text{'Forbes did it'} \rangle$, which is not. Since the latter is not a singleton, Rescher's Principle is not applicable. Applying Lewis's congruence definition, moreover, gives exactly the desired result: assuming the one element of this ordered set as given premiss raises the probability of the other; indeed it raises it to 1. The latter fact can even be taken in support of ascribing to Lewis acceptance of (3): full agreement is not just coherent; it is *very* coherent.

There is a complication that, although it needs to be addressed, does not affect the points just made. While the representation in terms of

¹⁰ Whether or not we have coherence at the level of 'assertions of supposed facts' will depend, as in the examples just given, on what assumptions are made concerning the reliability of the testimonies. But, as we hinted at in the beginning of this chapter and as will be clearer in the following chapters, the typical coherence applications involve full or partial ignorance as regards the reliability. Typically, then, we are not in a position to assess the coherence at the level of 'assertions of supposed facts'. If we were, that would mean that we had knowledge of facts of reliability, in which case we would have no use for coherence in the first place. The point is that there seem to be no interesting applications for a concept of coherence at that level.

simple ordered sets is sufficient to make sense of coherence as applied to testimonial agreement, there is still need for a minor amendment. Compare the case of Smith and Jones with another involving only Smith, who is, we suppose, queried on two different occasions, each time testifying to Forbes's being the culprit. This situation would, just like the Smith–Jones scenario, be represented by \langle 'Forbes did it', 'Forbes did it' \rangle , if the ordered-set policy is adhered to. But, unlike Smith's and Jones's agreeing with each other, Smith's agreeing with himself, albeit on different occasions, would normally not be a noteworthy fact, especially not if, as in this case, the agreement concerns a single simple proposition and not a long complicated story, the details of which may be hard to recall if they have been fabricated. The upshot is that we need to have a representation that allows us to distinguish between these two cases. A representation, accordingly, should provide the resources necessary for allowing us to determine the source of the information, in this case, the incriminating witness.

It turns out that the systems whose coherence is at issue are best thought of as sets of ordered pairs, each pair consisting of an assertion plus the proposition that is asserted. In the Forbes case, for instance, the relevant set is $\{\langle$ 'Smith says that Forbes did it', 'Forbes did it' \rangle , \langle 'Jones says that Forbes did it', 'Forbes did it' $\rangle\}$. By the coherence of this set we mean, as before, the coherence of the ordered set or sequence \langle 'Forbes did it', 'Forbes did it' \rangle . We are now better off, however, in the sense that we have the conceptual resources we need for determining the sources. The Smith–Smith example would be represented as $\{\langle$ 'Smith says at time t_1 that Forbes did it', 'Forbes did it' \rangle , \langle 'Smith says at time t_2 that Forbes did it', 'Forbes did it' $\rangle\}$, which is distinct from the set representing the Smith–Jones scenario.

This idea can be generalized along three different dimensions: by allowing (1) for more than just two supposed facts asserted, (2) for the supposed facts asserted to be different and not the same, and (3) for other sorts of evidence for a given supposed fact than evidence in the form of assertions. If we perform these three generalizations at once, what we end up with is the powerful concept of a *testimonial system* $\mathbf{S} = \{\langle E_1, A_1 \rangle, \dots, \langle E_n, A_n \rangle\}$, where E_i is any sentence of a report type providing putative evidence for supposed fact A_i . Such a testimonial system is, by definition, coherent just in case the ordered set of its supposed facts $\langle A_1, \dots, A_n \rangle$ is. My final proposal is to equate Lewis's

sets of supposed facts asserted with testimonial systems. The restriction to finite systems is no real limitation, since any given person can only receive a finite number of testimonies.

Returning now to Bonjour, if one looks at his theory from an abstract point of view, the systems he is concerned with, and to which he suggests the concept of coherence be applied, are, without exception, testimonial systems in my sense. First of all, he assumes that coherence can be assessed relative to the belief system of a given person. Important here is the Doxastic Presumption saying that the enquirer may take her believing this and that as *bona fide* facts, that is, as something that is not under dispute (1985: 101–6). Thus the person is supposed to have available ‘reports’ from her belief system that this or that is believed. Based on these reports and the coherence of their contents, she is roughly supposed to be able to assess the acceptability of the beliefs. What Bonjour is describing is, in effect, a testimonial system where the incoming reports are of the form ‘S believes that A’, i.e. the following sort of systems: $\{\langle \text{‘S believes that } A_1 \text{’}, A_1 \rangle, \dots, \langle \text{‘S believes that } A_n \text{’}, A_n \rangle\}$. Testimonial systems of this kind will be called *doxastic* systems. At the same time, Bonjour insists that the concept of coherence be applied also to cases of witness agreement. He remarks, in his comments on Lewis’s examples about the witnesses’ telling the same story,

as long as we are confident that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth-telling as the only available explanation of their agreement.

Saving the portion of this extract that is about coherence and truth for later, we just note that Bonjour associates witness agreement with ‘a high enough degree of coherence’. What is common to witness cases and the sceptical scenario is, of course, that they are both instances of testimonial systems.

Now given the passage just quoted, there can be little doubt that Bonjour would subscribe to (1), the meaningfulness of applying coherence to testimonial agreement. The extract can even be taken in support of (2) and (3): agreement, we are told, amounts to a ‘high enough degree of coherence’.

It remains to be assessed whether Bonjour’s multi-aspect theory of coherence also supports this interpretation. In view of the vagueness of

that theory, I have only a tentative argument to the effect that it does. In the light of the foregoing remarks on testimonial systems, I will read his five coherence criteria as being applicable not only to ‘systems of beliefs’, in terms of which they are explicitly cast, but to testimonial systems generally. Now if we take BonJour’s third criterion literally we should, when assessing the coherence of a given testimonial system, focus on two factors: the *number* of inferential connections between the elements and the *strength* of those connections. These factors should then somehow be amalgamated into one coherence judgement. In the case of perfect agreement, there are just two inferential connections due to the mutual implication between what the witnesses say. These connections, being logical in nature, are as strong as they could possibly be, meaning that perfect agreement fares extremely well as regards strength. On the other hand, it obviously does not do too well as regards the number of connections, there being, as noted, only two of those. But the idea of simply counting the connections seems, on second thought, naive. A large but scattered system may have more connections than a small but tightly interwoven set, a fact which hardly prevents us from regarding the smaller set to be the more coherent one. This suggests that the number of connections should be normalized somehow in order to reduce the dependence of coherence on the sheer size of the system. One way of achieving this would be to divide the number of *actual* connections with the number of *possible* ones. Relative to this amended measure, perfect agreement fares much better, indeed extremely well. The strength factor, as just noted, is already at its maximum. And in this case the number of actual connections equals the number of possible ones. So, not only is (2) satisfied on this more plausible reading of the third criterion; (3) is validated as well. Agreement is not only coherent; it is very—indeed maximally—coherent.¹¹

2.4 Truth and Agreement

If coherence theorists can agree on nothing else, they should at least grant that full agreement is a case of coherence, and perhaps of a high

¹¹ On this rendering of BonJour’s criteria, it does not matter, as far as the degree of coherence is concerned, how many witnesses attest to the same thing. For more on coherence and size, see sections 6.1 and 7.4.

or even maximum degree of coherence. The significance of this fact for our concerns lies in its opening up the possibility of at least a partial assessment of the otherwise notoriously unclear problem of coherence and truth. Could it at least be shown that agreement implies truth? While a positive answer to this question would lend considerable plausibility to the general thesis that coherence implies truth, further investigation of other forms of coherence would be needed to settle the matter. The situation would be quite different if the answer turned out to be negative. If agreement—the paradigm case of coherence—turns out not to imply truth, then this would amount to a convincing refutation of the general claim that coherence does. More carefully put, this eventuality would confront the coherence theorist with a dilemma: either she would have to concede that coherence in general does not imply truth or she would have to loosen the conceptual tie between agreement and coherence. Thus, either coherence does not imply truth or it is ‘robbed of all significance’.

Nonetheless, the relevance of the question whether agreement implies truth could be questioned. This is especially true if what we are interested in are primarily anti-sceptical uses of coherence. Central in that context is the notion of one person’s applying the concept of coherence to her own beliefs or memories so as to certify those beliefs or memories. But while full agreement makes perfect sense in the context of witness testimonies, it does not seem applicable at all to one person’s beliefs or memories. At a given moment in time, a person cannot believe or remember the same thing twice, and so no two of her beliefs or memories can be in perfect agreement. If two beliefs or memories have the same propositional content, these beliefs or memories are not distinct but one and the same. What gives rise to this peculiarity is not so much the fact that we are looking at beliefs or memories rather than, say, reports from witnesses. Rather, the trouble derives from the fact that in the sceptical scenario all beliefs or memories are supposed to belong to the same person and hence stem from one and the same source. There is no problem so long as the beliefs or memories are beliefs or memories of different persons. My belief that the sun is shining, for instance, may coincide with your distinct belief with the same content, making full agreement between our different beliefs possible.

My response to this objection is threefold. First, Lewis and Bonjour, two of the most prominent coherence theorists, are of a different opinion. For them agreement is a species of coherence, indeed it is a paradigm case of coherence, which is used to illustrate general claims about coherence and truth. Second, it turns out that a great majority of the remarks that I will make in connection with full agreement carry over to more general settings allowing less-than-full agreement between reports. The more general approach was investigated in Bovens and Olsson (2000). That study also indicates that several additional technical problems arise in the more general setting, without the corresponding pay-off in results that could not have been foreseen by studying the simpler case. It should be pointed out that Part II of this book examines the problem whether more coherence implies a higher likelihood of truth in a general set-up. Finally, the objection shows at best that full agreement cannot be realized in all types of testimonial systems. But this need not affect its status as an ideal form of coherence. The fact that there are no perfect circles in actual physical systems does not diminish the importance of the concept of perfect circularity as an ideal that real empirical circles can approximate to a greater or lesser degree. The same is true of full agreement which cannot be realized among one single person's beliefs or memories but which those beliefs or memories can come indefinitely close to realizing.

There is another competing intuition about coherence that needs to be addressed. Instead of taking witness agreement as a model of coherence, one can think of a coherent situation as similar to a jigsaw puzzle in which all pieces fit together so as to make up a meaningful picture. The reason why this intuition is in opposition to the framework adopted here is that the pieces that fit so well together in the jigsaw puzzle are qualitatively *different* pieces. From that perspective, perfect agreement, far from being a paradigm case of coherence, does not seem to make sense. What counts against the puzzle theory of coherence, however, is precisely the fact that it does not make possible a useful coherence assessment of full agreement among witness testimonies. Such cases in which, in addition, the reliability of the reports can be called into question are surely possible and even frequent. And the question may arise as to what can be said about the likelihood of truth of what is reported on the basis of the limited information at hand. In its blunt refusal even to make a coherence

assessment, the puzzle theory has nothing to offer in this respect. According to the agreement theory, on the other hand, it is meaningful to assign a degree of coherence in the hope of using it to estimate the likelihood of truth. I consider both intuitions worth exploring, though. In this book, I have chosen to follow C. I. Lewis, Laurence Bonjour, and, I believe, C. A. J. Coady in relying on the agreement intuition. Paul Thagard has presented an alternative theory based on the puzzle analogy. His theory is examined in section 9.4.

There is another objection to the claim that full agreement is a case of coherence. Consider a testimonial system $\mathbf{S} = \{\langle E_1, A_1 \rangle, \langle E_2, A_2 \rangle\}$ where A_1 and A_2 are two self-contradictory propositions. This would also be a case of full agreement, and yet the situation seems anything but coherent. Moreover, if this is a case of coherence, then coherence definitely does not imply truth. This objection can be seen to have little merit once the potential applications of coherence are brought into the picture. The whole idea, we recall, was to use coherence as a guide to truth. Surely we need such a guide only if we cannot decide the matter without it. If a reported proposition is internally contradictory, then we know that it cannot be true, and we know that any system to which it belongs cannot be true as a whole. Hence we have no need for coherence as a guide to truth. This shows that the real issue is whether coherence implies truth *for non-contradictory systems*. The coherence theorist claims that it does. She can claim this and at the same time consistently maintain that the system \mathbf{S} is coherent, although pathological because of the internal inconsistencies involved. For the purposes of the essay it may be assumed, without any loss of generality as far as possible applications are concerned, that the contents of individual reports are self-consistent.

2.5 A Simple Witness Model

Agreement among reports can establish a high probability of what is agreed upon, even though the individual reports, taken singly, would not be particularly good evidence for the proposition in question. If several witnesses agree point by point in their description of the presumed culprit, then we tend to think that what they say must be true. And we tend to think this even if we did not take those

witnesses to be terribly credible before they delivered their agreeing statements. I will now go on to underpin these intuitions using probability theory. I will then proceed to examine the conditions under which coherence has this impressive effect on the likelihood of truth.

I will use a simple probabilistic model of converging witnesses due to Michael Huemer (1997) to illustrate the effect of agreement. Huemer claims that the model adequately represents the sort of witness scenario that Lewis and Bonjour had in mind, a claim which we will find reason to question later (in Chapter 3 and 4), but the model is still useful as an illustration of some basic facts.

Let us return to the Forbes case. Smith and Jones, we recall, both incriminate Forbes, who is, we now assume, one among n possible culprits. Consider the following propositions:

H = 'Forbes did it'

E_1 = 'Smith says that Forbes did it'

E_2 = 'Jones says that Forbes did it'

We can calculate the posterior probability that Forbes did it, i.e. the probability of his guilt given the two testimonies, using Bayes's theorem:¹²

$$P(H/E_1, E_2) = \frac{P(E_1, E_2/H)P(H)}{P(E_1, E_2/H)P(H) + P(E_1, E_2/\neg H)P(\neg H)}$$

In order to calculate this we need some additional assumptions, the meaning of which will be elucidated in the subsequent sections. First of all, we will make two independence assumptions. In particular, we will assume that the reports are independent conditional on the truth as well as on the falsity of the hypothesis H , that is to say, $P(E_2/E_1, H) = P(E_2/H)$ and $P(E_2/E_1, \neg H) = P(E_2/\neg H)$. It follows that $P(E_1, E_2/H) = P(E_1/H)P(E_2/H)$ and $P(E_1, E_2/\neg H) = P(E_1/\neg H)P(E_2/\neg H)$. Second, we will suppose that Smith and Jones are equally credible individually. Letting i denote the reports' initial credibility, this amounts to assuming $i = P(E_1/H) = P(E_2/H)$. Our assumptions so far allow us to conclude $P(E_1, E_2/H) = P(E_1/H)P(E_2/H) = i^2$. Third, we will stipulate, reasonably, that if a given

¹² Unless otherwise stated, 'posterior probability' refers to the probability of the hypothesis conditional on the two testimonies.

testimony is false, any one of the $n - 1$ innocent suspects can be incriminated with equal probability. The probability that a given witness incriminates Forbes falsely is $P(E_1/\neg H) = P(E_2/\neg H) = (1 - i)/(n - 1)$.¹³ If we combine this with our assumption of independence conditional on the falsity of the hypothesis, what we get is $P(E_1, E_2/\neg H) = P(E_1/\neg H)P(E_2/\neg H) = (1 - i)^2/(n - 1)^2$. Collecting all this and plugging it into the original equation yields eventually

$$P(H/E_1, E_2) = \frac{ni^2 - i^2}{ni^2 - 2i + 1}$$

We note for later use that, in this model, the posterior joint probability depends on just two factors: (1) the number of equally likely suspects or, what comes to the same thing, the prior probability of what is agreed and (2) the individual credibility of each testimony.

To illustrate the effect of agreement, suppose there are initially ten equally likely suspects, so that $n = 10$. Since Forbes is assumed to be one of them, $P(H) = .1$. Suppose, further, that $i = P(E_1/H) = P(E_2/H) = .5$. Plugging this into the equation above yields $P(H/E_1, E_2) = .9$. Thus, taken in isolation neither testimony is sufficient to confer a high probability upon what is being said. And yet, when the two testimonies are combined, it becomes highly likely that what they say is indeed true. This is surely impressive considering the fact that we had just two agreeing reports.

2.6 Conditions for Convergence

If adding one more agreeing testimony increases the probability of what is being agreed upon, we will say that the testimonies *converge*. Formally, we have convergence if $P(H/E_1, E_2) > P(H/E_1)$. What has been shown thus far is that agreement *can* lead to convergence in this sense under suitable circumstances and that the increase in probability due to a second testimony can, in those cases, even be impressive. It is

¹³ The relevant probability is the chance that the witness gives a false testimony multiplied by the chance that he picks out Forbes falsely. A further reasonable assumption needs to be made in order for the former chance to equal $1 - i$. Let $E_{1,k}$ = 'Smith says that suspect k did it' and, similarly, $E_{2,k}$ = 'Jones says that suspect k did it'. Further, let H_k = 'Suspect k did it'. The extra assumption is $P(E_{1,k}/H_k) = P(E_{2,k}/H_k) = i$. For the chance that a given witness—say, Smith—gives a false testimony equals 1 minus the chance that he gives a true testimony: $1 - [P(E_{1,1}/H_1)P(H_1) + \dots + P(E_{1,n}/H_n)P(H_n)] = 1 - i$.

time to examine the nature of these favourable circumstances more carefully. Before we proceed to the more interesting conditions, we note that prior probability of the contents of the reports has been assumed to be non-extreme. Obviously, if a supposed fact initially has a probability of 0 or 1, no new evidence can change that assignment. This goes in particular for evidence in the form of one or more testimonies.

2.6.1 *Testimonial Independence*

One assumption that we made use of to illustrate the amplifying force of agreement was that of conditional independence. ‘Conditional’ here means ‘conditional on a specific assignment of truth-value to the hypothesis’. Two testimonies are conditionally independent just in case, once the truth-value of the hypothesis is known, what the one witness has said does not affect the probability of what the other witness will say. The assumption of conditional independence has two parts, corresponding to assuming the hypothesis true or assuming it false: $P(E_1/H) = P(E_1/H, E_2)$ and $P(E_1/\neg H) = P(E_1/\neg H, E_2)$. These two assumptions serve to simplify calculations tremendously and yet this is not their main motivation. Rather, there is a wide consensus among probability theorists that conditional independence captures, in probabilistic terms, the intuitive idea of *testimonial independence*, i.e. the idea that the reporters have not coordinated their testimonies in any way, for example by fudging their stories into agreement.¹⁴

What reasons are there for thinking that conditional independence in the sense just referred to is an adequate probabilistic representation of testimonial independence? We first note that such independence is not adequately represented by the equation $P(E_2/E_1) = P(E_2)$. According to this equation, Smith’s incriminating Forbes should make us neither more nor less confident that Jones will testify against Forbes as well. In many cases in which the witnesses have not in any way coordinated their testimonies beforehand this equation will not be satisfied. Given that Smith and Jones are intuitively independent and somewhat credible, Smith’s testimony against

¹⁴ See, e.g. Cohen (1977) and Jeffrey (1987).

Forbes will raise the probability that Jones will make the same assessment. The reason is that, Smith being somewhat credible, his testimony against Forbes will raise the probability that Forbes actually did it. Given Jones's partial credibility, the latter raise should in turn positively affect the probability that he will give a similar testimony.

The lesson to be drawn from the observation just made is that we must, in testing for real testimonial independence, make sure that such indirect dependences between testimonies have been blocked. This is accomplished by taking steps to ensure that the probability of Forbes's guilt is not affected by the assumption that one witness has delivered a positive testimony. This can be done either by setting the probability of Forbes's guilt to 1 (thus assuming Forbes guilty) or by setting that probability to 0 (thus assuming Forbes not guilty). Once the matter of Forbes's guilt has been settled, no further assumption can influence its probability. This holds in particular for assumptions about witness testimonies. The testimonies are independent, in the relevant sense, just in case there is no influence between them given that Forbes's guilt has been decided. This amounts to saying that they are conditionally independent in our sense of $P(E_1/H) = P(E_1/H, E_2)$ and $P(E_1/\neg H) = P(E_1/\neg H, E_2)$. In short, the characteristic feature of testimonial independence is that, while such testimonies may well be, and typically are, directly influenced by the supposed fact asserted and hence indirectly relevant to each other, there is no *direct* influence between the testimonies (Figure 2.1).

Is conditional independence necessary for convergence, i.e. for the addition of an agreeing testimony to have a positive effect on the




Figure 2.1: *Independent testimonies.* Smith's and Jones's testimonies are directly influenced by the fact they are reporting on. There is no direct influence between the testimonies themselves.

probability of what is being said? Intuitively, if the reporters are entirely dependent, so that the one is just repeating what the other has said without bothering about its truth or falsity, adding the second report should have no effect on the probability of what is agreed upon. We might as well choose to listen to just one of the reporters, disregarding the other. This intuition admits of simple probabilistic verification, as shown in Goldman (2001).

Although some independence is necessary for agreement to be significant, it is clear that the witnesses need not be completely independent for this to happen. L. Jonathan Cohen's well-known corroboration theorem establishes the truth of this preconception in probabilistic terms (1977: 101–7). As Cohen notes, it may be true of each of two positively relevant testimonies that neither corroborates the other. Thus, we might have (1) $P(H/E_1) > P(H)$ and (2) $P(H/E_2) > P(H)$ without having either (2C1) $P(H/E_2, E_1) > P(H/E_1)$ or (1C2) $P(H/E_1, E_2) > P(H/E_2)$. This raises the question under what circumstances the addition of an agreeing testimony does add to the support of what is agreed upon. According to the theorem, (2C1) is entailed by (1) and (2) supplemented by the conditions (3) $P(E_1, E_2) > 0$, (4) $P(E_2/E_1, H) \geq P(E_2/H)$, (5) $P(E_2/E_1, \neg H) \leq P(E_2/\neg H)$, and (6) $P(H/E_1) < 1$. With the additional condition (7) $P(H/E_2) < 1$, (1C2) follows as well. Thus, full conditional independence is not necessary; in the context of the other conditions it suffices to have the weaker conditions (4) and (5). That is to say, it suffices that there be no negative influences between the testimonies, if the hypothesis is true, and no positive influences among them, if the hypothesis is false.

Can we have convergence even if Cohen's conditions are not satisfied? Cohen appears to think that we cannot, for he claims that his conditions are 'severally necessary and jointly sufficient as a formal reconstruction under which corroboration takes place' (1982: 162–3). There can be no doubt that they are jointly sufficient. Nonetheless, as I will show in Chapter 3, they are not 'severally necessary'. The addition of another agreeing testimony may yield a positive effect even if Cohen's conditions do not hold. Whether it does so depends on the assumptions that are made in the specific case and there seems to be no general theorem like Cohen's to appeal to in such cases.

Another point is that although Cohen's theorem gives weak conditions under which the addition of an additional testimony makes a positive difference, it does not say how big that difference is. Presumably, the more independent the witnesses are, the bigger is that difference. I will not, however, attempt to verify this (very reasonable) conjecture but just note that such verification would require an account of 'degree of independence'.

2.6.2 *Individual Credibility*

As the reader can easily verify, another feature of Huemer's model is the individual credibility of the reports. Each report was assumed to be positively relevant to its content, however modest that relevance might be, so that $P(E_i/H) > P(E_i)$ which is, of course, equivalent to $P(H/E_i) > P(H)$.

We also note, for the record, that we assumed that the individual reporters are not *fully* credible, i.e. it was part of the example that $1 > P(H/E_i)$. The case of fully credible reporters is of no interest in this connection. If the reporters are initially fully reliable, one report is sufficient to certify the truth of what is being said. Since the posterior probability is already as high as it can ever get, there is no need to listen to further testimonies and no need to invoke coherence. An enquirer who is fortunate enough to have at his or her disposal fully reliable information sources has no use for coherence, the need for which arises only in the context of less than fully reliable information sources.

How essential is it to convergence that the reporters be individually credible? Huemer (1997) shows that there is, in his model, no convergence if the individual reporters are entirely useless. To see the point, we recall first that in Huemer's model

$$P(H/E_1, E_2) = \frac{ni^2 - i^2}{ni^2 - 2i + 1},$$

where n is the number of possible suspects and $i = P(E_1/H) = P(E_2/H)$. The reports have no individual credibility if $i = P(H) = 1/n$. Plugging this into the equation above yields $P(H/E_1, E_2) = 1/n = P(H)$. Hence, nothing is gained, as far as posterior probability is concerned, by combining independent reports that are individually useless: the result of such combination will be just as worthless as the original reports uncombined.

It is worth noting that this negative result relies on the other assumptions that are part of Huemer's model and that it does not just make use of the assumed individual uselessness of the reports. One of these additional assumptions is, of course, that of testimonial independence, but it was also assumed, for instance, that, if a testimony is mistaken, it is as likely to incriminate one given innocent suspect as any other.

Nonetheless, Huemer's observation is of utmost importance to the issue of coherence and truth since it shows that *full agreement, the paradigm case of coherence, does not have any effect on the likelihood of truth by itself*. Whether it has this effect depends on the circumstances and, in particular, on facts of credibility and independence.

2.7 Convergence Parameters

As we saw in section 2.5, the posterior joint probability of what is agreed upon depends (1) on the number n of equally probable suspects or equivalently the prior probability of the agreed proposition and (2) on the credibility i that each report has, taken in isolation.

In the simple scenarios we are now considering, the posterior probability of what is being asserted increases with its prior *ceteris paribus*, i.e. provided the credibility parameter i is held fixed. Suppose for instance that we fix i at $2/10$. Then setting $n=10$ yields $P(H/E_1, E_2) = 0.36$, whereas setting $n=100$ results in $P(H/E_1, E_2) = 0.861$. Later we will consider other witness scenarios where we do not have this simple relationship between the prior and the posterior. However, it will still be true that the latter depends on the former, albeit in more complex ways.

As one might have expected, the posterior probability increases with i , the credibility parameter. If the reports are more credible individually, the result of combining them will be a higher posterior probability *ceteris paribus*. Given $n=10$ and $i=1/10$, the posterior equals 0.1. Increasing the individual credibility to $i=2/10$ means raising the posterior to 0.264.

Although it suffices for our purposes to consider the two-report case, it is worth noting that the convergence parameters will also play a crucial role in how the probability of what is being agreed on

changes as further reports come in. The generalization of Huemer's model to the general finite case is straightforward (for a proof see Observation 2.1 in Appendix C):

$$P(H/E_1, \dots, E_w) = \frac{1}{1 + (n-1) \left(\frac{1-i}{i(n-1)} \right)^w}$$

As can be verified relative to this model, just how many reports it takes to reach a given probability sufficient for acceptance will depend on the prior probability of the agreed proposition as well as on the credibility of each report taken singly. A higher prior probability means that fewer reports will be needed to reach any one given level of posterior probability (*ceteris paribus*). The same goes *mutatis mutandis* for credibility: a lower individual credibility must be compensated for by an increase in the number of agreeing reports (*ceteris paribus*).

2.8 Challenges for the Coherence Theorist

I have illustrated some elementary facts concerning the potential for agreement to raise the likelihood of truth of supposed facts asserted. Furthermore, I have tried to elucidate the background conditions that were assumed to hold in the illustrating example. The two most interesting conditions were seen to be those of testimonial independence and individual credibility. The reports are independent if the reporters have not agreed beforehand to coordinate their reports by fudging them into agreement. I have indicated how the notion of testimonial independence can be spelled out uncontroversially in probabilistic terms. A report is individually credible, furthermore, if it is a somewhat, but not fully, reliable indicator of the truth of its content. As we also saw, the posterior joint probability in witness scenarios is dependent on two parameters: the prior probability of what is being agreed upon and the credibility of each report taken singly.

It should be emphasized that none of our observations so far has been established to hold for all possible witness scenarios. What has been said is true of what we might call Huemerian situations,

i.e. scenarios that can be represented using Huemer's model, but one would expect several of the observations to hold more generally. But, again, this has not been shown yet.

What has been shown conclusively, following Huemer, is that full agreement does not *by itself* imply a high likelihood of truth. To show this it was sufficient to display one single situation under which we have full agreement without having any boost at all in the likelihood of truth as the effect of this concurrence. We saw that this is exactly what happens in a Huemerian situation when the independent reports are individually useless.

But if full agreement—the paradigm case of coherence—does not by itself imply a high likelihood of truth, this would seem sufficient to disprove once and for all the general claim that coherence has this effect. And if coherence by itself does not imply a high likelihood of truth, the anti-sceptical use of coherence will be unsuccessful. The coherence theorist must show minimally that the conditions under which agreement is useless do not obtain 'in reality'. There are a number of possibilities here. One could hold, for instance, that our beliefs, memories, or whatever are individually credible a priori. We will encounter this contention in the works of Lewis and Coady. Another possibility would be to argue that Huemer's model is inadequate as a model of beliefs, memories, etc. This type of (reconstructed) response is actually implicit in all three theories that I will scrutinize in the rest of this part of the book: in Lewis's, in BonJour's, and in Coady's. As we will see, BonJour argues in addition that, once we focus on the epistemically relevant sort of scenarios, we can have an increase in the likelihood of truth as the result of concurrence even though each individual report is entirely useless. A further problem for these theorists is somehow to accommodate the presumed dependence of the likelihood of truth on the particular values of the convergence parameters.

C. I. Lewis's Radical Justification of Memory

3.1 The Problem of Justifying Memory

Can we trust our memories? What reasons, if any, do we have for believing that events we remember (in the non-veridical sense of 'remember') actually happened? Obviously, we may have a lot of evidence that a particular recollection represents an event that actually took place. However, on closer scrutiny that evidence will be seen to depend in turn on other recollections. The special philosophical problem about memory, or at least one problem that has received considerable attention, is that of justifying memory *as a faculty* without appealing to any recollections in the process.

The task of providing a radical justification of memory seems formidable. Presumably almost everything we believe, at least in the empirical domain, relies in some way or the other on memory, and yet a non-circular general justification of memory must not include a reference to any such recollections. Hence, it seems that we cannot make use of any of our empirical beliefs in the validation of memory. If I try to assess the credibility of my present empirical belief, then, in Lewis's own words,

at each step I shall be sent back to past experience; hence to the evidence of memory; and shall always find that the available evidence of memory is insufficient and requires itself a ground of credence, which in turn, can only be found in past experience—as remembered—and so on. The general nature of memorial knowledge constitutes a Gordian knot. (337)¹

But if we cannot appeal to our empirical beliefs, then what *can* we appeal to? In response to this problem Lewis suggests that a person

¹ As always, references to Lewis concern his book from 1946.

can validate his or her memories by examining the degree to which they cohere, for 'when the whole range of empirical beliefs is taken into account, all of them more or less dependent on memorial knowledge, we find that those which are most credible can be assured by their mutual support, or as we shall put it, by their *congruence*' (334). Such congruence, he maintains, raises the probability of what is remembered to the level of practical certainty in a way analogous to that in which agreement of testimonies can eventually make us convinced that what is being testified is true.

Yet, as we noted in Chapter 2, agreement—the paradigm case of coherence—apparently has the desired effect only under special circumstances that involve independence and individual credibility. At least this is strongly suggested by our simple witness model. Moreover, referring to the same model, even if those conditions are satisfied, a high enough probability is not secured unless the convergence parameters—the degree of individual credibility and the prior probability of what is being asserted—have been favourably fixed. Thus Lewis must, so it would seem, show not only that these conditions are satisfied in the case of memories, but also that the convergence parameters are determinable and that their values, once determined, are such as to guarantee a sufficiently high probability of what is seemingly remembered. The problem is that there seems to be nothing (or at least nothing empirical) to appeal to in the verification of these preconditions because, as we noted, preciously little can be taken for granted in a justification of memory from scratch. We are left with the impression that the appeal to congruence is ultimately unsuccessful as a reply to the sceptic. Granting that the effect of coherence can be impressive if the circumstances are favourable, the sceptic will insist that the anti-sceptic should justify her claim that such favourable conditions obtain in the actual case. In this chapter we will take a closer look at Lewis's proposal for how to solve these problems, of which Lewis himself shows admirable awareness.

3.2 Lewisian Witness Scenarios

Lewis takes the predicament of a person's wondering whether her memories are true at large to be analogous to that facing a juror who

is to decide what to make of the different testimonies presented to the court. In this section I will enquire into the nature of the sort of witness scenarios that Lewis had in mind. I will then argue, drawing on Horwich's probabilistic analysis of 'surprise', that the Lewisian set-up corresponds to the case in which agreement would be a surprising fact. Finally, I will show how this sort of scenario can be modelled in precise probabilistic terms.

3.2.1 *Lewis on Witness Corroboration*

In Chapter 2, we made acquaintance with Michael Huemer's probabilistic model of agreement (Huemer 1997). Huemer uses his model to underpin Lewis's claim that there are circumstances under which the congruence of reports establishes a high probability of what they agree upon, even though no single report would be particularly good evidence for anything in question. Huemer seems to think that this 'fairly modest' (ibid.: 468) statement exhausts Lewis's view on the relation between congruence and likelihood of truth. This, as we will see, is incorrect, since Lewis also advanced stronger theses that reach well beyond this item of common sense. The point I wish to make at this point, however, is that Huemer's model—which he refers to as 'the Lewisian example'—is inadequate as a representation of the sort of witness situation that Lewis actually contemplated (ibid.: 471).

Let us return to Lewis's 'relatively unreliable witnesses who independently tell the same story'. The relevant passage deserves to be quoted in full:

For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon, by principles of probability determination which are familiar: on any other hypothesis than that of truth-telling, this agreement is highly unlikely; the story any one false witness might tell being one out of so very large a number of equally possible choices. (It is comparable to the improbability that successive drawings of one marble out of a very large number will each result in the one white marble in the lot.) And the one hypothesis which itself is congruent with this agreement becomes thereby commensurably well established. (246)

That is to say, 'out of all the possible ways in which unreliable reporters can go wrong, their happening to tell independently just these stories which agree point for point, would be so thoroughly incredible on any other hypothesis than that of accurately telling the truth' (352). That which is congruent 'becomes thus probable just in measure as this congruence would be unlikely on any supposition which is plausible' (ibid.).

In the case highlighted by Lewis, there are initially two open hypotheses about the reliability of the witnesses: they are either 'truth-telling' and hence reliable or 'false', in which case they just generate their reports at random from a large number of equally possible alternatives. Before we have queried the reporters, we do not know which hypothesis to accept. However, upon observing the agreement, we become more inclined to think that the witnesses are telling the truth, as agreement would be highly unlikely on the alternative hypothesis of a mere random selection. As an effect of the increased probability of reliability, we also become more confident that what the reports say is true.

Needless to say, none of this should be taken to imply that there is a need to revise the *prior* probability of the reliability hypothesis because of congruence. A low antecedent probability of reliability is consistent with a high posterior probability of reliability. Given no evidence, the chance that the reports are reliable may be low, but once the different reports come in, the probability of reliability may, in the light of their agreement, become much higher than it was before.

The problem with Huemer's model, as a representation of a Lewis scenario, is that there is nothing in it that corresponds to the two alternative hypotheses that are supposed to be available, i.e. the hypotheses of reliability vs. unreliability. As a consequence, it cannot do justice to the intuition that observed agreement should make us more convinced that the witnesses are reliable.

One might expect of a Lewisian scenario that agreement on something relatively specific, and hence antecedently relatively improbable, should confer a relatively high probability upon the hypothesis that the witnesses are reliable. For instance, agreement on 'A one-legged, white-haired man with a beard committed the robbery' should make us more convinced that the witnesses are reliable

than should agreement on ‘A man committed the crime’. This is so because agreement on the first proposition is less probable than agreement on the second one given the alternative (random) hypothesis. Furthermore, as the probability of reliability increases, so should the probability of the proposition agreed upon, or so one might be inclined to think. Although these intuitions are on the right track, the exact relationship between specificity and the probability of the agreed proposition turns out to be a considerably more complicated matter. I will consider it in section 3.2.3 and return to it in Chapter 7 in connection with a discussion of L. Jonathan Cohen (see section 7.5).

3.2.2 *Lewis Scenarios and Surprising Agreement*

Paul Horwich (1982) has provided a general probabilistic account of the circumstances under which an event is, properly speaking, surprising. Having noted that the improbability of an event is a necessary but not a sufficient condition for the event to be surprising, Horwich goes on to specify what further conditions might distinguish improbable events that are surprising from those that are not. He observes that what probability we assign to an event derives from our opinions about the circumstances under which it occurred. Suppose, for instance, that I am about to toss a coin 100 times. The reason why I assign a low probability to 100 consecutive heads is my belief that the coin is probably fair. Let C represent the beliefs about the circumstances and E the statement that may or may not be surprising. For E to be surprising it must hold that E is initially highly unlikely, i.e. $P(E) \approx 0$. The further condition which Horwich proposes is $P(C/E) \ll P(C)$; that is, ‘the truth of E is surprising only if the supposed circumstances C , which made E seem improbable, are themselves substantially diminished in probability by the truth of E ’ (101). An application of Bayes’s theorem yields

$$P(C/E) = \frac{P(E/C)P(C)}{P(E)}$$

Let K represent some alternative account of the circumstance that we initially assign a low probability. Then $P(E) = P(E/C)P(C) + P(E/K)P(K) + P(E/\neg C, \neg K)P(\neg C, \neg K)$. As Horwich notes, the

requirement that $P(C/E) \ll P(C)$ will be satisfied if $P(E/K)P(K) \gg P(C)P(E/C)$, that is, 'if there is some initially implausible (but not wildly improbable) alternative view K about the circumstances, relative to which E would be highly probable' (102).²

Applying Horwich's analysis to the case of corroborating testimonies, the statement which may or may not be surprising is $E = E_1 \& E_2$, the fact that Smith's and Jones's testimonies agree on Forbes, whereas C may be a statement saying that the witnesses are reliable to a certain degree, e.g. that each is right in 60 per cent of the cases. It can also express the utter unreliability of Smith and Jones. Whatever C is, agreement is never surprising if we know, or think we know, that C is true; if the probability of C is 1, this assignment will not be changed as the effect of taking into account the new information provided by the testimonies, and so the condition $P(C/E_1, E_2) \ll P(C)$ will be violated. For agreement to count as surprising there must be an alternative hypothesis regarding the reliability, which must be, to some extent, unknown or uncertain.

Suppose now that the circumstances are such as described by Lewis, i.e. that we believe initially that the witnesses are most likely to be completely unreliable (i.e. no better than randomizers) but also allow for the possibility that they may be fully reliable. Our 'alternative view about the circumstances' will then be the hypothesis that Forbes did it and both Smith and Jones are telling the truth. If the prior probability of agreement is very low, and the alternative view initially implausible but not wildly improbable, then, since agreement on Forbes is highly probable given that view, Horwich's conditions are satisfied, and agreement counts as a genuinely surprising fact.

3.2.3 *Modelling a Lewis Scenario*

My next undertaking will be to describe Lewis's witness scenario in probabilistic terms. To confirm its adequacy in this respect, I will show that it allows for agreement on something relatively improbable

² Several authors have given essentially the same account of what makes something surprising or in special need of explanation. John Leslie expresses the idea as follows: 'A chief (or the only?) reason for thinking that something stands in [special] need [of explanation], i.e. for justifiable reluctance to dismiss it as how things just happen to be, is that one in fact glimpses some tidy way in which it might be explained' (1989: 10). For a similar view, see van Inwagen (1993: 135).

to make it relatively likely that the reports are reliable and, at a certain level of improbability, relatively likely that what is reported is true. In Olsson (2002*b*), I investigated a model of ‘uncertain reliability’ in the context of a discussion of some claims about the relation between the prior and the posterior due to L. Jonathan Cohen. Interestingly, the model employed there turns out to correspond exactly to the kind of set-up that Lewis found to be of such great importance.

Consider the following variation on the Forbes scenario. Suppose there were two groups of witnesses observing a criminal act, one group standing close to the crime scene and the other far away. Unfortunately, we do not know to which group a given witness belonged (and it is no use asking the witnesses because, we assume, they are all unreliable regarding their location at the relevant time). For all we know, a given witness may have been standing close to the crime scene or she may have been standing far away. Now, two witnesses, our old friends Smith and Jones, have been selected to identify the criminal, whereupon they independently incriminate Forbes (H).

In this situation there is uncertainty about the reliability of Smith and Jones. If they were standing close to the scene of the crime, they would be considered reliable. If, on the other hand, they were standing far away, they would rather be judged unreliable. In order to simplify the situation further we will assume that they would be completely reliable in the first case and completely unreliable in the second. We will stipulate, again in the interest of avoiding unnecessary complexity, that Smith and Jones were standing at the same place.³ We let R stand for ‘Smith and Jones are both completely reliable’ and U for ‘Smith and Jones are both completely unreliable’. R and U are mutually exclusive and exhaustive hypotheses about the reliability profiles of Smith and Jones, meaning that $P(R) + P(U) = 1$. The probability that Forbes did it, $P(H)$, will be assumed to be non-extreme.

If Smith and Jones were standing close to the crime scene and hence are fully reliable, then the following hold:

$$P(E_1/H,R) = 1 = P(E_2/H,R) \quad (\text{i})$$

$$P(E_1/\neg H,R) = 0 = P(E_2/\neg H,R) \quad (\text{ii})$$

³ This assumption will be subject to critical scrutiny below.

If Smith and Jones were standing close by and Forbes is guilty, they will both be sure to identify Forbes as the criminal (by (i)). Further, we may exclude the possibility of both making a mistake at such short a distance (by (ii)).

We recall that in order to obtain a situation in which agreement may be surprising, such agreement must be initially unlikely. We shall accomplish this by assuming, first, that unreliability is antecedently by far the more probable hypothesis about the circumstances and, second, by assuming that the probability of Smith's incriminating Forbes at random decreases with the probability of Forbes's guilt, so that it is relatively unlikely that Smith picks out Forbes at random if Forbes is relatively unlikely to be guilty. To make things simple we will stipulate that the probability of Smith's incriminating Forbes purely by chance *equals* the probability of Forbes's being guilty, and similarly for Jones.

$$P(E_1/H,U) = P(H) = P(E_2/H,U) \quad (\text{iii})$$

$$P(E_1/\neg H,U) = P(H) = P(E_2/\neg H,U)^4 \quad (\text{iv})$$

Under what circumstances would this be realistic? Let us imagine that Smith and Jones are presented with a line-up comprising all and only the suspects of the case, Forbes included, among which they have to choose, and that the suspects are equally likely to be the criminal in question. Then the probability that Forbes did it is $1/n$, where n is the number of suspects, and if Smith and Jones are completely unreliable, they will each pick out Forbes with probability $1/n$, regardless of whether Forbes is actually guilty or not.

We have made the substantial assumptions already; the rest is just independence. We assume that Smith and Jones pick out their suspect independently. This is trivially satisfied if they are reliable, since then $P(E_1/H,R) = 1 = P(E_1/H,R,E_2)$. What we need to stipulate is that they incriminate independently if they are unreliable:

$$P(E_1/H,U,E_2) = P(E_1/H,U) \quad (\text{v})$$

$$P(E_1/\neg H,U,E_2) = P(E_1/\neg H,U) \quad (\text{vi})$$

⁴ Conditions (iii) and (iv) together imply $P(E_1/H,U) = P(E_1/\neg H,U)$, that is to say, the hypothesis and the evidence are independent conditional on a witness's being unreliable. This condition has previously been imposed in the so-called 'Scandinavian School of Evidentiary Value'

Our next two independence assumptions stipulate that the issue of Forbes's guilt or innocence is of no consequence to the issue of Smith's and Jones's exact whereabouts at the time of the crime:

$$P(R/H) = P(R) = P(R/\neg H) \quad (\text{vii})$$

$$P(U/H) = P(U) = P(U/\neg H) \quad (\text{viii})$$

Thus, for instance, assuming Forbes guilty neither increases nor decreases the probability that Smith and Jones were located near the place where the crime was committed—surely a reasonable stipulation.

We are now in a position to compute the posterior probability of H on the evidence provided by E_1 and E_2 . As shown in Appendix C (Observation 3.1), that probability is given by

$$P(H/E_1, E_2) = \frac{P(R) + P(H)^2 P(U)}{P(R) + P(H)P(U)}$$

It can be verified that corroboration takes place in this model, i.e. we have $P(H/E_1, E_2) > P(H/E_1)$ provided $P(U)$ is non-extreme. That is to say, two testimonies are better than one (for a proof see Observation 3.4 in Appendix C).

Let us see how the posterior depends on the prior in this setting. There are two issues here: one is how the prior probability of the hypothesis affects its posterior probability; the other is how the prior probability of the hypothesis affects the posterior probability of reliability. As one can imagine, these issues are intimately connected. We will begin with the former question.

In order to make it initially unlikely that Smith and Jones are reliable, we assume that there were relatively few witnesses standing close to the scene of the crime, so that it is antecedently unlikely that Smith and Jones should belong to that distinguished group. Figure 3.1 shows the posterior probability of the hypothesis as a function of the prior for $P(U) = 9/10$ and $P(R) = 1/10$. The posterior joint

(Schum 1988: 258–61). For details, see Edman (1973), Hansson (1983: 78), and, for the connection with my model, Bovens, Fitelson, Hartmann, and Snyder (2002: 555–6). I was not aware of the history of the condition when I stated it in Olsson (2002*b*). If my endorsement of (iii) and (iv) makes me belong to the Scandinavian School, I would not complain about it, being a Scandinavian myself.




Figure 3.1: *The posterior probability of the hypothesis as a function of its prior.*

probability decreases with the prior down to a certain point. Once that point is reached, the posterior rises steeply as the prior is further diminished.

Why does the curve in Figure 3.1 look the way it does? Its shape can be explained in terms of how the prior probability of the hypothesis influences the probability that Smith and Jones were standing close to the crime scene and hence are reliable. Under our assumptions the latter probability can be calculated as follows:

$$P(R/E_1, E_2) = \frac{P(R)}{P(R) + P(U)P(H)}$$

R and U have been assumed to be exclusive and exhaustive possibilities, and so $P(U/E_1, E_2) = 1 - P(R/E_1, E_2)$. The posterior probabilities of R and U are plotted in Figure 3.2 as functions of the prior probability of the hypothesis. We can now explain the curve in Figure 3.1. Why is the posterior decreasing with the prior at first? If the prior is high, then unreliability is by far the most probable alternative (Figure 3.2) and so the reports will have no or little effect on the posterior joint probability, which will approximate the prior (Figure 3.1). Why, at a certain point, does the posterior probability of the hypothesis start increasing as the prior is further diminished? The probability of unreliability decreases




Figure 3.2: *The posterior probability of reliability vs. unreliability as a function of the prior probability of the hypothesis.*

with the prior probability of the hypothesis, and so the probability of the remaining alternative—reliability—increases (Figure 3.2). As it becomes increasingly likely that the witnesses are reliable, it also becomes increasingly likely that the hypothesis is true (Figure 3.1). This, after all, is what the witnesses keep telling us. As we noted before, this is exactly the kind of phenomenon we would expect from Lewis’s own description of the sort of case to which he assigns such epistemological significance.

It is easy to see that in a Lewis scenario the degree of individual credibility is determined once we know both the prior probability of what is being said and the prior probability of reliability. To be exact, $P(H/E) = P(R) + P(H)(1 - P(R))$ (for a proof see Observation 3.2 in Appendix C). It is natural to regard $P(H)$ and $P(R)$ as the two fundamental convergence parameters in a Lewis setting. This is what I will do in the following, although I will not always distinguish sharply between individual credibility and prior probability of reliability.

In a Lewisian scenario, as modelled here, nothing is gained by combining individually useless testimonies. There is no convergence if the testimonies have no individual credibility or, equivalently, if the probability of reliability is zero (see Observation 3.5 in Appendix C for a proof).

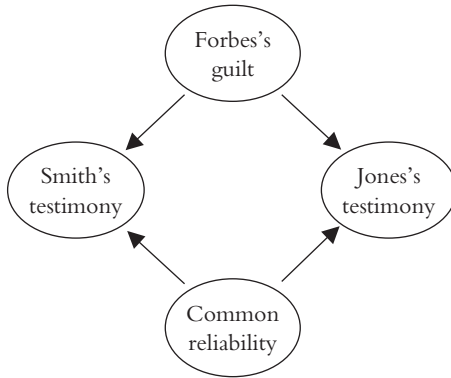


Figure 3.3: *Dependent reliability in my model.*

There is a complication that needs to be addressed at this point. We saw, in Chapter 2, that Smith's and Jones's testimonies are independent in the standard sense if and only if the following conditions hold:

$$P(E_2/H, E_1) = P(E_2/H)$$

$$P(E_2/\neg H, E_1) = P(E_2/\neg H)$$

As Bovens and his colleagues have shown, neither of these conditions is satisfied in my model.⁵ It is instructive to see why independence fails in this model. As Cohen points out, independence requires that the reports of the witnesses be causally connected only through the truth of what is being said. The problematic assumption in my model is the simplifying supposition that Smith and Jones are either both reliable or both unreliable. As Bovens et al. observe, this means that the reports are not only connected causally through the truth of what is being reported; they are also connected through their reliability profiles. This point is best conveyed by means of graphical illustration (Figure 3.3). What Bovens et al. prove, more specifically, is that the following hold in my set-up:

$$P(E_2/H, E_1) > P(E_2/H)$$

$$P(E_2/\neg H, E_1) > P(E_2/\neg H)$$

⁵ See Theorems 1 and 2 in Bovens, Fitelson, Hartmann, and Snyder (2002).

It is not difficult to see why these conditions hold in my model. As for the first, if we know that Smith has correctly incriminated Forbes, we may strongly suspect that he was standing close to the crime scene and hence is reliable. It was part of my model that if Smith is reliable, so is Jones. Hence, Jones will probably give a true testimony as well, that is to say, he will probably also incriminate Forbes, given that he did it. At least, the chance that he will is greater than it would have been had we only known that Forbes did it without knowing about Smith's testimony. As for the second condition, if we know that Smith has erroneously incriminated Forbes, we may strongly suspect that he was standing far from the crime scene and hence is unreliable. Hence, Jones, who was standing at the same place, is probably also unreliable and will probably give a false testimony as well. The chance that he will do so by incriminating Forbes is greater than it would have been had we only known that Forbes did not do it without knowing about Smith's testimony.

Bovens and his colleagues have improved upon my model by freeing it of the assumption that Smith and Jones are either both reliable or both unreliable, so as to make room for genuine independence which they accordingly stipulate. Their model can be graphically depicted as in Figure 3.4. We note that the only causal link between the reports is through the truth of the hypothesis, as required by genuine evidential independence.

Nonetheless, the shortcoming of my model with respect to independence turns out to be entirely innocent. Everything that was said above about my model holds also in the amended, but slightly more computationally complicated, setting. In particular, agreement on




Figure 3.4: *Independent reliability in the model of Bovens et al.*

something relatively improbable will make it relatively likely that the reporters are reliable and, given a certain level of improbability, relatively likely that what is reported is true.

There is a further reason for not dismissing my model. Although it violates independence and hence is dubious as a probabilistic representation of a Lewis scenario, it happens to have surprising implications for another issue, namely for the tenability of Cohen's claim that his conditions are not only sufficient but also necessary for convergence. This issue was raised in Chapter 2 in connection with the elucidation of the convergence conditions. We recall Cohen's theorem according to which $P(H/E_2, E_1) > P(H/E_1)$, or (2C1), is entailed by

- (1) $P(H/E_1) > P(H)$
- (2) $P(H/E_2) > P(H)$
- (3) $P(E_1, E_2) > 0$,
- (4) $P(E_2/E_1, H) \geq P(E_2/H)$,
- (5) $P(E_2/E_1, \neg H) \leq P(E_2/\neg H)$
- (6) $P(H/E_1) < 1$.

Whereas Cohen's condition (4) is satisfied in my model, his (5) is not. In fact, all conditions except (5) are satisfied, provided $P(R)$ is non-extreme. For a proof of (1) and, by symmetry, (2), see Observation 3.3 in Appendix C. Cohen's conditions are not all satisfied and yet, as we saw, we still have convergence in the sense of (2C1). My model thus disproves the first part of Cohen's contention that his conditions are 'severally necessary and jointly sufficient as a formal reconstruction under which corroboration takes place' (1982: 162–3). The conditions are indeed jointly sufficient, but they are not severally necessary.⁶

As my model illustrates, convergence is possible even if Cohen's weak independence conditions are not satisfied. However, whether we actually have convergence will depend on the particular

⁶ Based on Cohen (1977), O'Neill (1982) ascribes to Cohen the claim that his conditions are not only sufficient but also necessary for corroboration to take place. However, it is not evident what Cohen actually means by 'corroboration' in his 1982 article. In this difficult paper he seems to suggest that corroboration should not be defined as consisting in the mere satisfaction of (2C1), but as involving the satisfaction of (1), (2), (4), and (5) as well. So, Cohen could be taken to claim that his conditions, or at least a majority of them, are *analytically* necessary for corroboration. My argument above would have no force against this contention; it would follow from Cohen's view, thus understood, that the case

assumption made, and to the best of my knowledge there is no result of a general nature comparable to Cohen's theorem to appeal to in that case.

3.3 Lewis on the Convergence Conditions

Although a Lewis scenario differs from Huemer's model in a crucial respect, our previous observations about the latter concerning the convergence conditions and parameters are equally true of the former. We know that agreement boosts credibility given that Cohen's weak conditions are satisfied, the actual level reached being dependent on the prior probability and degree of individual credibility. Finally, if the reports are individually without credibility, nothing is gained by combining them; the whole will be just as useless as its parts.

Lewis is acutely aware, far more so than any other coherence theorist I have encountered, of how important it is to ensure the satisfaction of the convergence conditions for coherence justification to work. He notes that congruence reasoning applies to justification of empirical beliefs, most of which are in his view based on memory, 'so far as these are believed initially on grounds which are independent, and are not based on the same evidence, or the one of them believed merely because the other has already been accepted' (347). If, by means of contrast, a coherent set is 'fabricated out of whole cloth, the way a novelist writes a novel, or if it should be set-up as an elaborate hypothesis ad hoc by some theorist whose enthusiasm runs away with his judgment, such congruence would be no evidence of fact' (352). If it were such evidence, then 'unreliable reporters would be working in the interest of truth if they got together and fudged their stories into agreement' (ibid.), a notion that Lewis, quite rightly, dismisses as absurd. As for individual credibility, he makes the following declaration: 'If...there

I describe is not one of corroboration at all, since (5) is violated. In his discussion of Cohen's theorem, Carl G. Wagner points out that the problem of what conditions are necessary has not always been clearly stated, for 'by asking for conditions which, along with (1) and (2), entail (2C1), one precludes at the outset the attainment of a set of conditions necessary for (2C1), for neither (1) nor (2) is necessary for (2C1)' (Wagner 1991: 297). Whether or not a watertight case can be made against Cohen on this particular issue, it is surely of interest to see that we can have (2C1) satisfied in the absence of independence in Cohen's very weak sense.

were *no* initial presumption attaching to the mnemically presented; no valid assumption of a real connection with past experience; then no extent of congruity with other such items would give rise to any eventual credibility' (357).

Lewis must argue that the conditions of independence and individual credibility are satisfied in the case of memory. This would allow him to appeal to the general fact that those conditions give rise to convergence (in the context of some other conditions of lesser importance). These conditions, furthermore, must be shown to be acceptable as true from a 'sceptical' point of view that does not involve any substantial empirical beliefs that rely on memory for their justification. Unfortunately, Lewis does not clearly separate the independence and credibility issues, addressing explicitly only the latter. This is clearly a most serious omission in the case of memories, for which independence is a particularly problematic assumption. As all students of witness psychology know, an eyewitness's individual report from memory may be influenced, not only by what actually took place, but also by her own prior attitudes and expectations. If, for example, she is firmly convinced that most thieves are foreigners, she will be more likely to 'remember' the thief as 'foreign-looking'. A further problem for the reader is that Lewis advances two rather different arguments for individual credibility, without clearly distinguishing them. One can be classified as a 'transcendental argument' aiming to show that the individual credibility of memory is a condition for the possibility of our 'sense of reality'. The second line of reasoning is verificationist in nature. It attempts to establish that 'systematic deception' is not a genuine possibility distinct from 'reliability' because a victim of such a deception would never be in a position to verify that she is one.

3.3.1 *The Transcendental Argument for Individual Credibility*

What is to be proved is that 'mnemic presentation itself should, before any further examination as to coherence, afford some probability of past fact' (356). Lewis's fundamental thesis is that this is something that does not admit of genuine doubt (357–8):

[T]his assumption in question has a certain kind of justification, in the fact that mnemic perseveration of past experience; its present-as-pastness; is

constitutive of the world we live in. It represents that continuing sense of reality beyond the narrow confines of the merely sensibly presented; the only reality which as humans we can envisage; the only reality which could come before us to be recognized as such. If we adopt the Cartesian method of doubting everything that admits of doubt, we must stop short of doubting this. Because to doubt our sense of past experience as founded in actuality, would be to lose any criterion by which either the doubt itself or what is doubted could be corroborated; and to erase altogether the distinction between empirical fact and fantasy. In that sense, we have no rational alternative but to presume that anything sensed as past is just a little more probable than that which is incompatible with what is remembered and that with respect to which memory is blank.

What Lewis is offering here is a transcendental argument for the initial credibility of memory. When we dismiss a contention as ‘mere fantasy’, we do so on the basis of what we perceive to be empirical facts. But, in Lewis’s view, what we take to be empirical facts ultimately rely for their justification on what we remember to be the case. If we started to doubt the initial credibility of memory, we would accordingly lose any criterion by means of which we could distinguish fact from fantasy. If we, as rational agents, want to uphold this distinction, we need minimally to acknowledge that the fact of remembering is positively relevant to that which is remembered. Thus, the individual credibility of memory is a condition for the possibility of the fact–fiction distinction.

A. J. C. Coady, whose theory of natural testimony will be considered in detail in Chapter 5, makes some remarks on memory that are astonishingly close to Lewis’s view on the matter. In particular, he offers the following illustrative example (1992: 98). Consider an individual who discovers that his memories are false because they clash with his sensory experience. The man thinks he remembers a large flowering gum tree at a certain place in a familiar park but when he goes there to admire it he cannot find it, though he soon comes across it in a nearby golf course which he recalls frequenting. Coady now remarks:

Any suggestion, however, that the general reliability of one’s memory is to be established by present perceptual experience cannot be seriously entertained since, unless we take its reliability for granted to some extent, we cannot even gather the empirical evidence which is supposed to make

the case for or against memory's connection with reality. The position of the tree may have been misremembered but to establish this we have to accept at face value a large number of memory deliverances, such as, that *this* is the park in question, *that* the golf course frequented in the past, and *this* the previously encountered tree, not to mention the fact that present observation itself rapidly assumes the status of memory . . . (ibid.: 98).

Coady offers the tree example as an objection to what he takes to be Hume's view, that the reliability of memory is discovered by experience, but the general point is identical with Lewis's contention that without some credibility attached to memory as such we would 'lose any criterion by which either the doubt itself or what is doubted could be corroborated'.

Lewis's thesis is that the *prima facie* credibility of memory is constitutive of our sense of empirical reality (361). Lewis regards this contention as akin in spirit to Kant's thesis about the transcendental unity of apperception. Since the assumption of a *prima facie* credibility of memory is constitutive of our sense of empirical fact, it is in no need of justification, the genuine difficulty being not to justify it but to formulate it (361).

Let us grant Lewis's point that individual credibility is a condition for the possibility of the fact–fiction distinction, so that we must accept the truth of this condition in order to make it. Obviously this will not convince a sceptic who is unwilling to treat the tenability of the fact–fiction distinction as a fundamental assumption that cannot be rationally doubted. What is so important about the fact–fiction distinction that makes it rationally impossible to give up? Lewis's first answer to this question is that without it our words would lack any meaning, making human communication impossible. In his own words, an initial credibility of memory is 'constitutive of our sense of the only reality by reference to which empirical judgments could have either truth or falsity or any meaning at all' (361). Without it, 'there could be no such things as fact and no intelligible discourse' (ibid.). In a final passage, he goes one step further in arguing for the pragmatic indispensability of the fact–fiction distinction:

we have no alternative but to accept the principle that mnemonic presentation constitutes a *prima facie* probability of past actuality, and to accept, in some form or other, the Rule of Induction. These are ingredients in, and together with the certainties of given experience, are constitutive of, *our sense of that*

reality which we cannot fail to acknowledge, unless we would repudiate all thought and action and every significance of living. That we cannot do. (362, my italics)

The reasons Lewis offers here for the fact–fiction distinction are partly pragmatic in nature, referring to ‘action’ and the ‘significance of living’.

The most impressive of these arguments is surely the one that focuses on meaning and intelligible discourse. For the sceptic, as she is understood in the Cartesian tradition, does not deny that words have meanings and that she is engaged in intelligible discourse with the non-sceptic. If it could be made likely that these presupposed conditions entail a fact–fiction distinction and that this distinction in turn requires the initial credibility of memory, that would seem to be a powerful argument against scepticism. It is an unfortunate fact that Lewis’s remarks in this direction are so impressionistic and unsystematic. Yet he should be credited for having provided what seems to be the first transcendental argument for individual credibility. We will later encounter similar arguments in the works of Coady and Davidson in connection with natural testimony and beliefs, respectively.

Lewis, as we just saw, tries to justify the fact–fiction distinction with reference to its importance for thought and action. In Chapter 10, I will argue that such pragmatic considerations should be brought in earlier in the argumentative chain. On closer scrutiny, we lack an incentive to reconsider our actual commitment to the high reliability and even trustworthiness of memory on neutral ground. From this perspective, Lewis emerges as a half-hearted pragmatist who failed to invoke his pragmatic principles wherever they are applicable.

3.3.2 *The Verificationist Argument for Individual Credibility*

Let us now turn to the other argument offered by Lewis. Although Lewis, in his transcendental argument, concludes that it is rationally impossible to engage in Cartesian doubt as to the credibility of memory and the associated fact–fiction distinction, he decides nonetheless to try carrying out such doubt, admitting that this must be, as he puts it, ‘an essay in the fantastic’ (358). As a first step, he invites us to suppose that a particular knower is subject to a systematic delusion of memory concerning his experience of music. As the person remembers it, whenever he has heard music in the past, it has

been accompanied by kaleidoscopic patterns of imaged colour. As a preliminary we are asked not to doubt the reliability of inductive generalization from the mnemically presented. The question Lewis now poses is whether such a knower will later discover, through the test of experience, the delusive character of this class of his memories. Lewis answers in the negative. Let us see why.

Suppose that music is promised, so that the person predicts the accompanying kaleidoscopic patterns, but that those patterns are not forthcoming. Given that the delusion does not extend to given sense experience, this event will be perceived as a puzzling exception to all past experience as the person remembers it. He will take note of the exception and decide to bear it in mind and, based on that memory, downgrade the reliability of similar predictions in the future. And yet, the next time music is promised, he will remember his past experience erroneously and predict, no less confidently, that the colour patterns will appear.

The general point Lewis wishes to make is that, while a victim of systematic delusion will continually note the necessity of revising downward the credibility of her recollections of certain types, she will always remember incorrectly the cognitive disappointments that motivated that revision in the first place, believing instead that her prediction was correct. The net effect is that she will never actually start doubting the credibility of those recollections. Thus, a person who is subject to a systemic delusion of memory will not be in a position to know this for a fact:

Our lives, as some outside omniscient observer would view them, are going to be a continual succession of disappointments of our cognitive expectations by our presented sense experience, but we are never going to know that general fact. Each disappointment is going to be strictly temporary, like the suffering some people suspect that they have endured under an anaesthetic, though the memory of it—they suppose—is later blocked. Also, what is a *cognitive* disappointment, may in other respects be a pleasant surprise. And on any grounds, occasional cognitive disappointments must be expected by anyone who commits himself to predictions that are less than certain. In short, we are going to lead quite normal lives. (359–60)

Based on his examination of the possibility of a systematic delusion, Lewis draws the general conclusion that '[t]he world as revealed to us

by our sense of past experience must be the world we live in' (360). For, as his example indicates, 'the distinction between an objectively real world and a sufficiently systematic delusion of one—congruent throughout—is a distinction which makes no discoverable difference, and is not the subject of any reasonable discussion' (*ibid.*). Lewis is here invoking the verification theory of meaning, in the tradition of Peirce, according to which a distinction that does not make any sensible difference is no real distinction at all. Needless to say, few contemporary authors would be willing to follow Lewis on this path.

One could also object that Lewis has so far only examined one way in which our memories could fall short of individual credibility, namely due to a systematic delusion, neglecting the further possibility of a general randomness. To this objection, Lewis replies that 'a general *unreliability* of memory would be something quite different and must reveal itself, if we preserved our rationality, in its falling into incongruities' (360). The suggestion is that we could detect a general unreliability (randomness) of our recollections by examining their degree of congruence. If they are highly incongruent, we may suspect unreliability. This is a puzzling statement in the light of Lewis's earlier observation that the real challenge is one of showing that our memories are somewhat credible *prior* to considerations of congruence. The reader is left with the impression that Lewis has failed to eliminate the possibility of complete unreliability and, as a consequence, that his justification for assuming individual credibility is unsuccessful.

Nonetheless, our observations in connection with Lewis's witness model indicate that this shortcoming in Lewis's argumentation need not affect his overall project. In a Lewis scenario, as I have represented it, it is not necessary to eliminate complete unreliability in order to secure the individual credibility of each report. It suffices to show that the reports are either completely reliable (*R*) or completely unreliable (*U*), and that neither of these possibilities has zero probability. Given that Lewis has succeeded in eliminating Cartesian deception, something that of course may be contested, a case could be made that the remaining possibilities are precisely *R* and *U*, and no further argument would be needed to ensure individual credibility. Lewis may have been taking on a heavier burden of proof than he had to.

The argument just considered relies on the verification theory of meaning and will therefore fail to convince most contemporary

philosophers. There is also evidence suggesting that Lewis considered this piece of reasoning, which he describes as an essay in the fantastic, to be less fundamental than his transcendental argument. I take it, then, that his basic view was that the initial credibility of memory is something that we cannot rationally give up, since doing so would mean undermining the fact–fiction distinction and, with it, the possibility of linguistic meaning and rational discourse. In the next section, I will argue that Lewis's attempted vindication of memory fails even if we grant that he managed to secure the individual credibility of memories.

3.4 The Individual Credibility 'need not be Assigned'

Lewis now claims that, once the convergence conditions have been proved to hold, we will not only have convergence in the sense of obtaining higher probability as the result of combining several reports; given that the degree of congruence is high enough, the resulting likelihood of truth will even be sufficient for 'rational and practical reliance'. But this is a severely problematic statement if one believes, with Lewis, that full agreement is a paradigm case of coherence. For, as we have observed, the likelihood of truth can show great variation depending on the values of the convergence parameters, even though witnesses agree fully on a given proposition. Referring back to Figure 3.1, we can get a posterior almost as low as 0.5 and almost as high as 1 merely by varying the prior probability of the proposition agreed upon while keeping the probability of reliability positive. On the analogy between witness reports and reports from memory, we should expect any given degree of congruence (whatever 'degree of congruence' means more precisely) to be compatible with a rather low eventual probability of the contents of the memories, even if the individual credibility is positive. To be in a position to determine whether the likelihood of truth is high enough for rational and practical reliance we need additional information about the prior and also about the probability of reliability. It is not sufficient merely to know that the latter is positive.

If we apply the point just made to the special case of memories, it follows that the mere individual credibility of our memory does not

entail any substantial conclusions about how likely it is that our memories are true. And yet what Lewis has achieved thus far is, at best, only that our memories have some, however small, positive individual credibility. Nothing in his attempted justification of the individual credibility of memories indicates a *particular* degree of positive initial credibility as pertaining to our memories as such. Aware of this difficulty, Lewis argues that it points to a fundamental limitation in our knowledge (356–7):

It does not appear that we could, candidly, assign any particular degree to [the initial credibility of memory]. We seldom take cognizance of this initial presumption, because generalizations as to particular *classes* of our memories intervene between it and any matter to be attested by memory. That recollections of the recent past are comparatively reliable; of the remote past, unreliable; that our memory for faces and for what we have said is trustworthy, but our remembering of names and dates is not: such generalizations will be the proximate grounds on which the credibility of particular memories is assessed. But these are, of course, generalizations from past experience (of remembering, and of later confirming or disconfirming) and as such are presently available only in the form of remembered experience, and require for their own authentication the presumption of initial credibility of the merely remembered as such. And the degree of this initial credibility, we have said, is hardly assignable.

Lewis thus concedes that the degree of individual credibility of our memories cannot be determined. All we can say is that it is positive. Based on our previous remarks, this observation of a fundamental lack of cognitive access to the values of the convergence parameters is actually sufficient to establish the impossibility of Lewis's vindication project.

Regrettably, Lewis, who is otherwise so perceptive in spotting problematic assumptions, does not conclude the failure of his enterprise but is even quite optimistic as to its prospects:

the degree of this initial credibility, we have said, is hardly assignable. But it does not need to be assigned. A larger or a smaller such initial probability would have no appreciable effect upon the eventually determinable probabilities in question *beyond that of a difference in the extent of congruity with other mnemonic items and with sense presentation which could be required for building up eventual probabilities sufficient for rational and practical reliance.* (357, my italics)

Lewis is here contending, surprisingly, that the exact degree of initial credibility pertaining to memory 'does not need to be assigned'. This degree, he maintains, does not matter for the eventually determinable probabilities (i.e. the posterior probability of what is mnemically represented), apart from the fact that a smaller initial credibility would require a higher extent of congruity for those probabilities to reach a level sufficient for 'rational and practical reliance'. Lewis is downplaying the latter correct observation as if it were a matter of no particular importance. In reality, however, it is of utmost significance to his endeavours. For if we do not know what the degree of initial credibility is—and we have Lewis's word that we cannot know this even in principle—then how should we ever be able to determine whether the actual extent of congruity is high enough for rational and practical reliance? Suppose Peterson finds himself having this and that memory, wondering whether he can rely on them at large, rationally and practically. According to Lewis, Peterson may assume that whatever the degree of individual credibility may be, so long as it is positive there is always a degree of congruence sufficient to boost probabilities up to a level where the contents of the memories can be accepted as true. But learning this is of no help to Peterson if, as Lewis would insist, he has no clue what the actual degree of individual credibility is. Purged of this knowledge, he will not be able to determine whether the actual degree of congruence of his memories is high enough to warrant acceptance of their contents.

To summarize, although Lewis's attempt to establish the individual credibility of memory contains several intriguing proposals, especially the transcendental argument, some of its elements are severely problematic. What seems particularly questionable is his verificationist dismissal of liar hypotheses at the outset. But even if we grant that Lewis succeeded in establishing the individual credibility of our memories, he ultimately failed to respond appropriately to the problem of the convergence parameters. The moral he drew from his correct observation of a dependence of the posterior on the particular degree of individual credibility is inadequate. The posterior probability of what is agreed upon or remembered is seriously under-determined by mere facts of agreement or, more generally, by mere facts of coherence. If we assume, with Lewis, that we do not, and indeed cannot, know the value of the convergence

parameters in the case of memory, it follows that we cannot say anything of substance about the magnitude of the posterior. Our conclusion must be that Lewis failed to show that coherence, when applied to the totality of our memories, implies a likelihood of truth high enough to warrant rational acceptance.

3.5 A Note on Lewis's Definition of Independence

I would like to make one final remark on Lewis's definition of independence and his thesis that coherence does not have any effect unless the individual cohering sources are in some degree credible, taken singly. The matter is important since it shows that in defining a Lewis scenario as we did, we were, in one respect, rather charitable to Lewis. The point I am going to make also illustrates the subtle nature of testimonial convergence.

As to the thesis, he writes: 'If, however, there were no initial presumption attaching to the mnemically presented; no valid supposition of a real connection with past experience; then no extent of congruity with other such items would give rise to any eventual credibility' (357). The discussion so far suggests that this thesis is correct, for it holds under the assumption that the reports are testimonially independent *in the standard conditional sense*.

Interestingly, Lewis favoured an explication of independence making that notion distinct from full conditional independence. He introduces his concept when discussing the corroboration of a hypothesis through its testable consequences. Since the testable consequences of a hypothesis are (putative) indicators of its truth and hence are reports, in the abstract sense, on that hypothesis, his remarks on independence are easily generalized beyond the hypothesis-consequences setting.⁷

Now Lewis considers two propositions, *A* and *B*, to be independent consequences of a hypothesis *H* just in case *A* and *B* are 'so related that supposing *H* false, the finding of one of them true would not

⁷ Lewis explains the relevant notion of consequence as follows (1946: 343, notation adapted): 'when a hypothesis *H* is said to have consequences *C*, what typically is meant is that *H*, together with other statements which may reasonably be assumed, gives a high probability of *C*.' He adds, in footnote 5 on

increase the probability of the other' (344). He restates this idea several times, with little variation, writing for instance, 'in general, the consequences of a hypothesis are independent only in the sense that the establishment of one does not increase the probability of another on the assumption that the hypothesis is *false*' (349 n. 6, original emphasis). It was hence clear to Lewis that the relevant notion of independence is of a conditional nature, but he insisted on taking into account only independence statements conditional on the falsity of the hypothesis, leaving out independence statements conditional on its truth. While Lewis was on the right track, the kind of independence he arrived at is ultimately too weak to capture what we intuitively mean by full testimonial independence.

Ironically, Lewis's notion of independence is too weak to support his own contention that aggregating independent but completely unreliable reports does not do a thing for the probability of what is being said. This can be shown relative to a slight modification of Huemer's model. According to Lewis's conception, independence of two pieces of evidence only requires their probabilistic independence given the falsity of the conclusion and not given its truth. Exploiting the liberal nature of Lewis's account, we stipulate that $P(E_2/E_1, H) = 1$. In other words, if Smith correctly reports Forbes to be the culprit, then so will Jones. This implies that $P(E_1, E_2/H) = P(E_1/H)P(E_2/E_1, H) = P(E_1/H)$. Lewis-independence, moreover, requires that the following hold true: $P(E_1, E_2/\neg H) = P(E_1/\neg H)P(E_2/E_1, \neg H) = P(E_1/\neg H)P(E_2/\neg H)$. As before $P(E_1/H) = P(E_2/H) = i$ and $P(E_2/\neg H) = (1 - i)/(n - 1)$. Bayes's theorem now yields:

$$P(H/E_1, E_2) = \frac{ni - i}{i^2 + ni - 3i + 1}$$

The issue at hand concerns complete unreliability. Setting $i = 1/n$ results in

$$P(H/E_1, E_2) = \frac{n}{2n - 1}$$

which is greater than $1/n = P(H)$, so long as $n > 1$. Hence, the combined reports support the hypothesis, even though taken singly they

the same page, that '[i]t is not even essential that such a probability be high', that is to say, 'in general the same principle will apply wherever C is more probable than not'.

do not. If $n = 2$, for instance, the posterior probability of H is $2/3$ whereas the prior equals $1/2$.

Admittedly the argument just given is not entirely decisive, as it involves Huemer's simple model rather than a Lewis-type scenario. Nonetheless, it strongly indicates an internal problem in Lewis's view. It also illustrates the perhaps surprising general point that in the absence of (full) conditional independence individually useless data may prove collectively useful. There is, on the other hand, no guarantee they will, as there seems to be no general result like Cohen's to appeal to in those circumstances.

Laurence Bonjour's Radical Justification of Belief

4.1 The Problem of Justifying Beliefs

Laurence Bonjour has provided a complex and sophisticated coherence theory of empirical knowledge. I do not aspire here to give anything like a complete account of all aspects of that theory. Rather, I will be focusing on Bonjour's claims about the relation between coherence and truth, and on the question whether there is, as Bonjour contends, a probabilistic foundation for those claims. I would like to emphasize that it is not improper, but actually quite fitting, to study Bonjour's theory from a probabilistic point of view, considering the important role played by such reasoning in his 1985 book, where he makes frequent use of, among other things, Bayes's theorem which he assumes that he can rely upon in his anti-sceptical enterprise (1985: 181).¹

Bonjour's theory is structurally very similar to Lewis's in that it is based, in considerable degree, on the witness analogy. Just as agreement between witness statements makes what is agreed likely to be true, so coherence among beliefs is supposed to make the contents of those beliefs amenable for rational acceptance. What correspond to the witness reports in Bonjour's case are reports to the effect that a person *S* believes this or that. Bonjour assumes that while *S* may not know that the content of a given belief is true, *S* does know that he has a belief with the given content, or at least we may take it for granted in epistemological enquiry that he knows this (1985: 81). That is what Bonjour's Doxastic Presumption states:

[T]he essential starting point for epistemological investigation is the presumption that the believer has a certain specific belief, the issue being

¹ Unless otherwise indicated, references to Bonjour are to his 1985 book.

whether or not the belief thus presumed to exist is justified, but the very existence of the belief being taken for granted in the context of the epistemological inquiry. And the further suggestion is that this presumption—that the believer in question does indeed accept the belief in question—though clearly empirical in content, is for these reasons available as a premise, or at least can function as a premise, in this context without itself requiring justification.

Thus, reports of the type ‘S believes this or that’ play essentially the same role in Bonjour’s theory as do reports of the kind ‘S seems to remember this or that’ in Lewis’s framework.

The idea now is, roughly, that a person is allowed to trust his acknowledged beliefs provided they show a sufficient degree of coherence. However, there are several complications that make this basic suggestion seem untenable, as witnessed by the many traditional objections to it. One of these objections is the so-called input objection. A pure coherence theory would seem to allow that a system of beliefs be justified in spite of being utterly out of contact with the world it purports to describe, so long as it is, to a sufficient extent, coherent. But this is surely an absurd result. A system of beliefs cannot be justified unless it receives some sort of input from outside and is thus causally influenced by the world.

BonJour admits that the force of the input objection must, to some extent, be conceded and hence that the purest sort of coherence theory—the one which does not allow for input from the world—must be deemed unacceptable. Nonetheless, he does not see this concession as a reason for endorsing foundationalism, insisting that ‘a theory which is recognizably coherentist—and more importantly free of any significant foundationalist ingredients—can allow for such input’ (110). This is a surprising move since it would seem that the subject engaged in radical justification of beliefs is, from her own sceptical perspective, unable to distinguish those beliefs that are observational from those that are not. We recall that what she is supposed to take for granted, in the sceptical exercise, are only facts of a report kind, that this or that is believed. The most such a subject can say, then, is that she *believes* this or that belief to be caused by outside forces.

As I understand Bonjour he is willing to concede that the subject does not have access to any facts concerning what beliefs are

externally caused and which are not. So, the coherence theory cannot admit input in the sense of 'externally caused beliefs'. Bonjour's point is that there is nonetheless a class of beliefs that can provide the subject with a sort of ersatz input. Membership in this class, moreover, can be assessed from the subject's internal perspective. This class is that of 'cognitively spontaneous beliefs' consisting of those beliefs that are acquired non-inferentially.

Cognitive spontaneity is intimately connected to observation, for observational beliefs are adopted non-inferentially. Suppose I sit at my desk. When I look around I come to have many beliefs as the result of observation: that in front of me there is a computer; that beside the computer there is a printer, and so on. It would be implausible to say that I have inferred that there is a computer and that there is a printer beside it, or that these beliefs were acquired through some other process of conscious deliberation. These putative facts strike me or occur to me in a way that is automatic and involuntary.² While cognitively spontaneous beliefs provide the subject with a kind of input, it is clear that their role cannot be to act as a foundational source of justification. The constraints on a pure coherence theory do not allow the cognitively spontaneous beliefs, or anything else for that matter, to have a degree of justification in themselves. Rather, to the extent that they are justified, it must be possible to argue for their truth from the believer's own perspective.

So far it has been established at best that cognitively spontaneous beliefs can provide the subject with cognitively accessible ersatz input and that those beliefs, just like any others, are to be justified with reference to other beliefs. Yet no requirement has been made to the effect that a system, in order to qualify as justified, must actually contain cognitively spontaneous beliefs. For all Bonjour has said so far,

² But are the cognitively spontaneous beliefs really accessible from the sceptical standpoint? We recall that from that point of view only statements of a report character may be taken for granted as true, i.e. statements of the form 'I believe this and that'. Thus the subject must not take for granted that a given belief was acquired in a cognitively spontaneous manner, but only that she *believes* that it was thus acquired. Against this, Bonjour might argue that whether a given belief arose spontaneously or not is a fact of consciousness with respect to which the subject may be assumed infallible. There are indications, however, that he would not want to follow that track—above all his employing the expression '*apparently* cognitively spontaneous beliefs' (148, my italics), which strongly suggests fallibility on the part of the subject. But if facts of spontaneity cannot be taken for granted, then they will have to be justified by facts of coherence. But this leads to a vicious circle since facts of the former kind were supposed to provide the coherence machinery with its input. I will ignore this difficulty for Bonjour in the following.

a system could be justified, in virtue of its high degree of coherence, and yet receive no input in the form of cognitively spontaneous beliefs. Bonjour proposes, accordingly, that this possibility be excluded:

[W]ithout input of some sort any agreement which happened to exist between the cognitive system and the world could only be accidental and hence not something which one could have any good reason to expect. Thus . . . a coherence theory of empirical justification must require that in order for the beliefs of a cognitive system to be even candidates for empirical justification, that system must contain laws attributing a high degree of reliability to a reasonable variety of cognitively spontaneous beliefs. (141)

In Bonjour's view, that there is a need for this Observation Requirement is an a priori truth but whether or not the requirement is satisfied in a given case cannot be determined a priori (142). For any given system it is an empirical question to be decided 'purely on the basis of coherence' whether the system, to the extent required, attributes reliability to some members of the general class of cognitively spontaneous beliefs (*ibid.*).

We are now in a position to express the sense in which Bonjour takes coherence to be connected with likelihood of truth:

A system of beliefs which (a) remains coherent (and stable) over the long run and (b) continues to satisfy the Observation Requirement is likely, to a degree which is proportional to the degree of coherence (and stability) and the longness of the run, to correspond closely to independent reality. (171)

This is a much more sophisticated claim than the one we began with, which stated plainly that coherence implies truth. For one, Bonjour is here making use of the notion of a 'degree of coherence' and not simply of coherence as a matter of all or nothing. His claim, furthermore, is contingent on the satisfaction of the Observation Requirement and on the (degree of) stability of the belief system 'over the long run'.

Bonjour's thesis is not only more sophisticated; it is also correspondingly more difficult to put to a test. In order to do so, we would have to assign definite meaning to the central concepts: 'degree of coherence', 'degree of stability', and 'likeness to correspond to independent reality'. As for degrees of coherence, there is, as we saw in Chapter 2, not much of substance to be derived from Bonjour's own theorizing. More will be said in Part II about degrees of coherence and likeness of truth as predicated of a system of beliefs.

The stability aspect, in contrast, will be largely ignored in this book. This is a justified omission given the purposes of this study. While Bonjour is ultimately aiming at defeating scepticism, our aim here is primarily to shed light on the supposed relation between coherence and truth. For us, but not for Bonjour, scepticism is a secondary issue, although it remains one of the most important applications of a probabilistic study of coherence. My discussion of Bonjour in this chapter will be concerned with those parts of his 1985 book where he makes comparatively clear statements about coherence and truth. I am here referring to his discussion of Lewis's witness example in his chapter 7 and those parts of his chapter 8 that deal with the possibility of deception.

In a central passage of his book, Bonjour notes that his coherentism appears to collapse into a Lewisian weak foundationalism with the cognitively spontaneous beliefs playing the part of Lewis's individually credible mnemonic presentations. But this, he hastens to add, is mere appearance. The difference is that '[a]ccording to such a foundationalist view, it is true *prior* to any appeal to coherence that cognitively spontaneous beliefs have this minimal degree of credibility—for which no adequate justification is or ever can be offered' (143). For a coherence theory, on the other hand, '*all* epistemic justification of empirical beliefs depends on coherence' (*ibid.*). This is a clear statement that, for Bonjour, cognitively spontaneous beliefs can have a minimal degree of credibility only subsequent to and not prior to coherence considerations, and so it would seem that Bonjour's theory is not a species of (weak) foundationalism after all.³

³ At another place, Bonjour gives the following characterization of the difference between coherentism and weak foundationalism (143): '[T]here is no reason why the Observation Requirement should be regarded as in fact conferring such an initial degree of warrant on cognitively spontaneous beliefs. The main point is that it is quite consistent with the Observation Requirement... that no cognitively spontaneous belief of any kind might turn out to be warranted: this would be so, for example, if no class of such beliefs turned out to be in internal agreement to any significant degree. But such a result would not seem to be possible for a weak foundationalist view, according to which the largest consistent (or coherent?) class of basic beliefs will seemingly have to be justified to some degree, even if perhaps not enough to satisfy the requirement for knowledge.' Here the contrast is made, not in terms of the probability of a given cognitively spontaneous belief *prior* to any appeal to coherence, but in terms of the probability of that belief *subsequent* to such an appeal. It is maintained that a coherence theory, but not a weak foundationalism, is compatible with the latter probability not exceeding the prior probability of the proposition in question. This, at least, is one way of reading this somewhat obscure passage.

The point just made takes us to the heart of the matter. Bonjour is apparently proposing that coherence can have an effect on probabilities of belief contents even if those beliefs are lacking in individual credibility. In his own words, ‘it is simply not necessary in order for such a [coherentist] view to yield justification to suppose that cognitively spontaneous beliefs have some degree of initial or independent credibility’ (147). It is to this central claim I now turn.

4.2 Bonjour on Justification from Scratch

Bonjour’s claim about the possibility of coherence justification ‘from scratch’ seems to contradict what we have been able to observe thus far. Both in the simple model and in Lewis’s model the witness reports were supposed to be individually credible, i.e. credible before any appeal to coherence is made. Only then was it possible to observe a positive effect of agreement on the posterior joint probability. Nothing, it seems, could be accomplished by combining independent but individually useless testimonies.

In his attempt to justify his surprising contention, Bonjour argues that C. I. Lewis’s example of ‘relatively unreliable witnesses who independently tell the same circumstantial story’ falls short of establishing that an individual positive credibility is required for agreement to boost probability. As we recall, Lewis argued that agreement of reports establishes a high probability of what they agree upon, since on any other hypothesis than that of truth-telling such agreement would be highly unlikely. This was supposed to occur although ‘[f]or any one of these reporters, taken singly, the extent to which it confirms what is reported may be slight’. Bonjour adds:

What Lewis does not see, however, is that his own example shows quite convincingly that no antecedent degree of warrant or credibility is required. For as long as we are confident that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth-telling as the only available explanation of their agreement—even, indeed, if those individual reports initially have a high degree of negative credibility, that is, are more likely to be false than true (for example, in the case where all of the witnesses are known to be habitual liars). And by the same token, so long

as apparently cognitively spontaneous beliefs are genuinely independent of each other, their agreement will eventually generate credibility, without the need for an initial degree of warrant. (148)⁴

We note the proviso 'so long as apparently cognitively spontaneous beliefs are genuinely independent of each other'. I am not aware of any argument on Bonjour's part that they are. I will return to the issue of independence in section 4.4. My concern here will be with individual credibility.⁵

Huemer (1997) uses the fact that in his model lack of individual credibility implies lack of convergence to counter Bonjour's thesis. While Huemer should be credited for seeing the crucial relevance of this issue to the sort of coherence theory that Bonjour advocates, his simple model is inadequate as a representation of the kind of witness scenario that Bonjour is concerned with. Bonjour's claim is made in connection with his discussion of Lewis, and we have seen that Huemer's model does not adequately represent a Lewis scenario. In the latter type of case, there are initially two competing hypotheses about the reliability of the witnesses: they are either truth-tellers or no better than randomizers. The proof in Chapter 3 about the impossibility of justification from scratch in a Lewisian set-up is therefore more decisive against Bonjour than Huemer's argument.

⁴ Note how Bonjour defines 'negative credibility' here. Reports have negative credibility, we are told, if they 'are more likely to be false than true'. Report *E*, then, is negatively credible in Bonjour's sense just in case the probability of *H*, given *E*, is less than one half: $P(H/E) < 0.5$. Compare this with the standard definition in terms of $P(H/E) < P(H)$. The two definitions come to the same thing if $P(H) = 0.5$. But if $P(H) < 0.5$, which is the case if, for instance, we have more than two equally likely suspects, negative credibility in Bonjour's sense is compatible with standard positive relevance, which is absurd. Take a criminal case in which there are ten possible equally likely suspects so that the probability that a particular one of them, Forbes, is the real culprit (*H*) is $1/10$. Suppose further that the positive testimony of Smith is capable of increasing the probability of *H* to $2/5$, in which case Smith's testimony, although positively relevant in the standard sense, has negative credibility in Bonjour's sense. Furthermore, given Bonjour's understanding of negative credibility, 'no credibility' presumably boils down to there being a fifty-fifty chance that *H* is true given a report to that effect, i.e. $P(H/E) = 0.5$. By the same token, this corresponds to the standard definition of irrelevance in the case of two equally likely suspects. If there are more such suspects, however, these two conceptions will diverge and what has zero credibility in Bonjour's sense may have a positive credibility in the standard sense. I will adopt a charitable interpretation of Bonjour in the following, assuming that he does not really want to dissociate himself from the standard definition of negative relevance (and irrelevance). The passage just quoted which suggests otherwise will accordingly be regarded a mere slip of the pen.

⁵ Note also that Bonjour seems to presuppose that we are able to identify the contents of beliefs independently of the issue of their individual credibility. This assumption could be questioned on the grounds that, if beliefs were totally unreliable, we would not be able to identify their contents. We will return to this 'Davidsonian' theme in connection with Coady (Chapter 5) who makes essentially this point in connection with natural testimony. Davidson's own theory will be discussed in brief terms in section 9.2.

Yet there are hints in the passage just quoted that the witness scenario BonJour has in mind may not be exactly identical with a Lewis set-up. What suggests this is BonJour's talk of 'habitual liars'. In a Lewis setting, again, there are supposed to be only two live possibilities as concerns the reliability profiles of the reporters, namely reliability or randomization. Lewis wanted to exclude lying at the outset, arguing that it could be dismissed on verificationist grounds before any appeals to coherence are made.

With or without liars, the following proof is general enough to handle all scenarios involving independent witnesses and should settle the matter quite conclusively to BonJour's disadvantage. Suppose, then, that our witnesses have no credibility whatsoever, that is to say,

$$P(E_i/H) = P(E_i)$$

Assuming that H is neither true nor false with certainty so that $0 < P(H) < 1$, (1) entails that $P(E_1/H) = P(E_1)$ and $P(E_2/H) = P(E_2)$, and also $P(E_1/\neg H) = P(E_1)$ and $P(E_2/\neg H) = P(E_2)$. Assume further that E_1 and E_2 are independent reports in the testimonial sense, so that they are probabilistically independent given that the truth-value of H is known. Formally,

$$P(E_1, E_2/H) = P(E_1/H)P(E_2/H)$$

and

$$P(E_1, E_2/\neg H) = P(E_1/\neg H)P(E_2/\neg H)$$

It is now straightforward to show that $P(H/E_1, E_2)$ equals $P(H)$ (see Observation 4.1 in Appendix C). Thus, if neither of two independent pieces of evidence is relevant to H , their concurrence has no impact on H 's probability.

Unlike Huemer's proof and unlike the proof presented in Chapter 3 in the context of Lewis's theory, the present demonstration does not rely on any additional assumptions beyond that of testimonial independence. Assuming independence is quite appropriate here given that our target is BonJour's contention that 'so long as apparently cognitively spontaneous beliefs are genuinely independent of each other, their agreement will eventually generate credibility, without the need for an initial degree of warrant'. Our derivation shows this claim to be in blatant error. What is true is that, so long as

apparently cognitively spontaneous beliefs are genuinely independent of each other, their agreement will *not* eventually generate credibility unless there is an initial degree of warrant.

BonJour's central argument for his coherence theory is manifestly false. Contrary to what he thinks, coherence cannot generate credibility from scratch when applied to independent data. Some reports must have a degree of credibility that is prior to any consideration of coherence, or such agreement will fail to have any effect whatsoever on the probability of what is reported. This holds, in particular, for cognitively spontaneous beliefs.

I would like to emphasize the generality of the foregoing simple demonstration. It applies to the simple Huemer model as well as to the more complicated Lewisian model: as soon as our independent reports lack individual credibility, we will not experience any rise in probability as the effect of agreement and the posterior will equal the prior. But the proof has still wider range, applying also to possible witness scenarios distinct from those envisaged by Lewis and Huemer. For example, it goes through whether or not 'lying' is considered an initial possibility alongside 'randomizing' and 'truth-telling'. This shows that admitting lying as seriously possible does not affect the present issue. Nonetheless, bringing in liars does have some other interesting consequences that will be considered next.

4.3 Lying and Individual Credibility

The attitude towards sceptical hypotheses reflects a genuine difference between BonJour and Lewis, who argued, on verificationist grounds, that 'the distinction between an objectively real world and a sufficiently systematic delusion of one—congruent throughout—is a distinction which makes no discoverable difference, and is not the subject of any reasonable discussion' (Lewis 1946: 360–1). This position, we recall, led Lewis to exclude radical delusion from consideration before any considerations as to coherence. In BonJour's view, by contrast, 'it will not do simply to close one's eyes [to sceptical hypotheses] by "refusing to entertain the skeptical question"' (180). While he is aware that '[s]ome attempt to argue, on broad verificationist grounds, that the skeptical views in question are either

meaningless or else somehow not genuinely distinct from nonskeptical views' (ibid.), this is a position in which Bonjour himself declares he 'can see no merit' (ibid.). Bonjour's characterization of verificationism fits Lewis's view perfectly, even though Bonjour does not mention Lewis in this connection.

There are some remaining issues concerning lying that still need to be clarified. We have already quoted Bonjour's statement that 'as long as we are confident that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth-telling as the only available explanation of their agreement—even, indeed, if those individual reports initially have a high degree of negative credibility, that is, are more likely to be false than true (for example, in the case where all of the witnesses are known to be habitual liars)'. First of all, it is simply incorrect to say that independent agreement among reporters *known* to be habitual liars will 'eventually dictate the hypothesis of truth-telling as the only available explanation of their agreement'. If we really *know* at the outset that the reporters are liars, no amount of agreement will change this fact.

Let us be charitable and construe Bonjour in the way I suggested before, i.e. as claiming that our epistemic predicament is analogous to a witness scenario that is just like a Lewis scenario, except that the liar hypothesis is also regarded as an open possibility alongside 'randomization' and 'truth-telling'. Even so, there is a further issue that needs to be addressed. On the current reading, Bonjour is suggesting that admitting lying as seriously possible will automatically confer zero or even negative credibility on each individual report so that we have at best $P(H/E) = P(H)$. Is that really so? As Tomoji Shogenji (2002) has shown, this holds only under special circumstances characterized by there being either (i) only one way to lie where lying is initially as likely as truth-telling or (ii) several ways to lie where lying is initially more likely than truth-telling.

Let us get clear about what lying means. First of all, a liar invariably produces false reports. Yet while there is only one way to tell the truth, there are normally many different ways to lie. If there are n suspects, there are $n - 1$ ways to lie about who is the real culprit. We will assume that a liar is as likely to produce one of the possible false reports as any other. Hence, if there are n suspects, the probability

that a liar incriminated one of the innocent ones will be $1/(n-1)$. Probabilistically speaking, the liar hypothesis (L) can be expressed as follows:

$$P(E/H, L) = 0 \quad \text{and} \quad P(E/\neg H, L) = 1/(n-1)$$

Reliability, we recall, is characterized by $P(E/H, R) = 1$ and $P(E/\neg H, R) = 0$, and unreliability by $P(E/H, U) = P(E/\neg H, U) = 1/n$. We will assume that any given witness is either telling the truth, randomizing, or lying: $P(R) + P(U) + P(L) = 1$.

In consonance with our probabilistic rendering of a Lewis scenario, we will stipulate that the reliability profile of the witness is not in any way dependent on the truth of the main hypothesis:

$$P(R/H) = P(R), P(U/H) = P(U), \quad \text{and} \quad P(L/H) = P(L)$$

Now suppose that report E lacks individual credibility, so that $P(H/E) = P(H)$. It can be verified that this assumption entails, when $n > 2$, $P(L) = (n-1)P(R)$ and hence $P(L) > P(R)$; and when $n = 2$, $P(L) = P(R)$. Moreover, if $P(L) = P(R)$ and $n > 2$, then $P(H/E) > P(H)$. For proofs, see Observation 4.2 in Appendix C.

What is the import of this result? Well, we remember that Bonjour's main application is his attempted radical justification of belief, in which case our imagined initial position presumably is one of ignorance as to whether our information is reliable or not. In the absence of a better way of representing ignorance probabilistically, we seem obliged to assign to each possibility the same probability. This is tantamount to assigning each of R , U , and L a prior probability of $1/3$. Moreover, given an initial ignorant state, there seems to be no reason to restrict the number of possible contents a given cognitively spontaneous belief may have to 2. When I discover you in the restaurant, you are but one of many people who could have been sitting at that particular table, and if some other person had been sitting there instead, the cognitively spontaneous belief that I would have acquired as the result of my observation would have had a different content. In general, the content of a given cognitively spontaneous belief is just one among many possible contents and if my cognitively spontaneous beliefs are 'lying' they may be doing so in a great number of possible ways. Now what we have shown is that invoking these two assumptions—'uniform prior over the possible

reliability profiles' and 'more than two possible report contents'—automatically confers a positive degree of credibility on each individual report. It would seem difficult to avoid the conclusion that each cognitively spontaneous belief is to some degree credible even before any appeals to coherence have been made. But this conclusion contradicts BonJour's contention that cognitively spontaneous beliefs are initially lacking in credibility.⁶

The incompatibility is serious because it involves a fundamental assumption in BonJour's epistemology. After all, he takes as the hallmark of his coherence theory that it does not require given data to be individually credible; this is the very feature that is supposed to distinguish his theory from Lewis's weak foundationalism. The underlying error is to think that the issue of individual credibility can be discussed separately from the issue of what reliability profiles (truth-telling, randomization, and so on) are possible.

4.4 Coordinated Lying and Independence

In his criticism of C. I. Lewis's weak foundationalism, BonJour commits himself to an analogy between our initial epistemic predicament and the witness scenario he sketches, a scenario that coincides with a Lewis scenario except in being slightly more liberal so as to admit the initial possibility of lying. One feature of such a BonJour witness set-up is independence. The possible liars are supposed to be, not conspiring liars, but 'habitual' liars. Unless they are habitually conspiring, which seems far-fetched, this seems to suggest that they simply lie without colluding. Yet in his final discussion of coherence and truth, BonJour argues that the main alternative to the reliability of our cognitively

⁶ See Chapter 7 for a more detailed examination of the consequences of invoking the Principle of Indifference to represent ignorance. As I will point out, there are certain conceptual problems involved in such invocation, as we may get, in certain respects, different results depending on how, exactly, the principle is put to use. Nonetheless, what is at stake here is the individual credibility of the data, a feature that turns out not to be heavily dependent on the particular employment of the Principle of Indifference. But I do not want to say that invoking this principle is unproblematic in this case. The objection I level here against BonJour should be seen as revealing a dilemma: either we invoke the Principle of Indifference, and get individual credibility on the cheap, or we do not invoke it, in which case no probabilistic analysis of the kind carried out by BonJour would seem possible, as we would lack a probabilistic representation of our supposed ignorance regarding the reliability of our cognitively spontaneous beliefs.

spontaneous beliefs is not one of uncoordinated, but of coordinated lying. This leads him to consider various sceptical hypotheses, mainly the Cartesian demon 'who employs all his powers to deceive me into believing that there is an ordinary world of the sort that I think there is, even though nothing of the kind actually exists' (179).

The distinction between collusion and habitual lying is obviously of no consequence for the issue of individual credibility. However, it does raise concerns as regards independence. We recall the relevant independence conditions:

$$(i) P(E_1/H, E_2) = P(E_1/H)$$

$$(ii) P(E_1/\neg H, E_2) = P(E_1/\neg H)$$

Cohen, we recall, showed that convergence is in fact guaranteed under the following weaker independence conditions (given some other conditions):

$$(iii) P(E_1/H, E_2) \geq P(E_1/H)$$

$$(iv) P(E_1/\neg H, E_2) \leq P(E_1/\neg H)$$

Clearly, if we *know* that the reports are coordinated lies, we also know that the reports are not independent. More precisely, (iv) (and, by implication, (ii)) will not be satisfied: if any one reporter lies in a specific way, we know that any other will lie, too, and indeed in the same way. However, our epistemic predicament, as envisaged by Bonjour, is not one in which it can be assumed as known that we are the victims of a coordinated deception. Rather, such deception is initially merely one among several open possibilities. As we have seen, Bonjour also thinks that cognitively spontaneous beliefs are, or at least can be, independent. But do these two components of his view really harmonize? Is it consistent with independence to admit collusion, be it as a mere possibility?

The answer to this question is twofold. First, (i) and (iii) are approximately satisfied regardless of whether the liars collude or not; and, second, whereas (ii) and hence also (iv) are approximately satisfied for uncoordinated liars, they are violated for coordinated liars given that n is large, that is to say, given a large number of possible reports. See Observations 4.3 and 4.4 in Appendix C for proofs.

The last observation is the most interesting one. As we have noted, independence is trivially lost if we know that the reporters have engaged in coordinating their lying. What the last observation says is that the same holds, albeit to a lesser extent, if we follow Bonjour's recommendation and view coordinated lying as a mere serious possibility (alongside truth-telling and randomization), provided that there are many possible report contents. As we have noted, there is every reason in the world not to assume that the number of possible contents of a given cognitively spontaneous belief should be small. Hence, transferred to Bonjour's anti-sceptical discussion, the last observation entails that the cognitively spontaneous beliefs will not be genuinely independent. Compare this with Bonjour's declaration that 'so long as apparently cognitively spontaneous beliefs are genuinely independent of each other, their agreement will eventually generate credibility, without the need for any initial degree of warrant' (148). Bonjour is here implying, conversationally if not logically, that cognitively spontaneous beliefs can be independent, which as we just saw is not true. What we have arrived at, then, is another serious clash between the components of Bonjour's theory. The cause of the problem is the false assumption that the question of independence can be answered separately from the question of what reliability profiles are possible.

4.5 Consequences for Bonjour's Anti-scepticism

We have established, in my view conclusively, that justification from scratch given independent data is impossible. If independent reports are individually useless, nothing will be gained by combining them, however high the resulting degree of coherence may be. We have discovered, in addition, two serious tensions, if not outright contradictions, in Bonjour's anti-sceptical theory. These are clashes between, on the one hand, Bonjour's criticism of Lewis aimed at drawing a sharp line between a coherence theory and weak foundationalism and, on the other hand, his account of which reliability profiles are to be regarded as seriously possible concerning our cognitively spontaneous beliefs.

The first conflict concerns his statement, in connection with Lewis, that cognitively spontaneous beliefs have no individual credibility,

a claim that is contradicted by his account of seriously possible reliability profiles (given two extremely plausible assumptions). The other dissonance arises from the fact that BonJour claims, again in the context of his Lewis critique, that cognitively spontaneous beliefs are, or at least can be, independent. But this is not correct if, as he declares later in the book, collusion is to be regarded as a serious possibility alongside reliability and randomization.

What could BonJour do to extricate himself from this predicament? The thesis about the possibility of justification from scratch should be rejected once and for all so that BonJour's theory inevitably collapses into weak foundationalism. Once that thesis has been dismissed, there is no reason for BonJour to cling to the presumed lack of initial credibility of cognitively spontaneous beliefs, and so the first tension is thereby automatically avoided. Concerning the independence problem, there are two possibilities for a revision, neither of which is especially attractive from BonJour's point of view. He may choose to give up the notion of cognitively spontaneous beliefs being independent in the testimonial sense. The problem with this move is that it precludes reference to general results like Cohen's theorem in support of the alleged convergence of cognitively spontaneous beliefs. We would have no reason to believe that massive coherence among such beliefs should have any effect at all on their joint likelihood of truth, let alone that the effect should be sufficient to warrant rational acceptance. The other strategy would be for BonJour to make yet another concession to Lewis, beyond having to admit the initial credibility of cognitively spontaneous beliefs, and exclude collusion as a serious possibility before any appeals to coherence are made. Perhaps he could motivate this move by invoking some form of Lewisian verificationism. The problem, of course, is that verificationism is generally an unattractive position in which BonJour has declared he can see no merit.

Perhaps there are other ways to avoid the two tensions I have pointed to. Perhaps one can, for example, model 'habitual lying' in some other way which does not require that a habitual liar always lies. Furthermore, I have treated lying as one possibility. But could one not treat each way of lying (each lie that could be told) as a distinct possibility? Or why not add the possibility of anti-collusion? Anti-collusion occurs when witnesses confer with each other and

deliberately set out to deliver disagreeing reports. My basic aim here has been to show that there are some subtle incompatibility problems lurking here, of which BonJour shows no awareness. I have illustrated these problems using a fairly straightforward probabilistic representation of the epistemological assumptions. The burden of proof is now on BonJour's defenders to show that there is no incompatibility.

In fairness to BonJour, it should be mentioned that he has recently abandoned the project of justifying beliefs from scratch through coherence reasoning.⁷ I will return to his reasons for this move in Chapter 6. Interestingly, those reasons bear little resemblance to the objections levelled here, but have to do with certain difficulties pertaining to comparative assessments of coherence.

⁷ See BonJour (1999).

C. A. J. Coady's Radical Justification of Natural Testimony

5.1 The Problem of Justifying Natural Testimony

Our reliance on other people's testimonies is no less extensive than our reliance on our own memories. It does not take much reflection to appreciate that our knowledge would be but a fraction of what we ordinarily think it to be were it confined to that which we can verify for ourselves via perception and reasoning. Without trust in testimony, written and spoken, we would not even know our passport numbers, our parentage, dates of birth, most geographical facts, and so on. In ordinary life we go about relying on the word of others so long as there is no positive reason to question the reliability of the informer. In this fashion we are able to extend our belief systems in ways that greatly transcend our own limited perceptual horizons.

Just as we can ask what reasons we have for trusting our memories, we can enquire similarly into the underlying rationale for the reliance upon the word of others. What is the philosophical basis of this reliance? Can it be vindicated in some way by showing that the testimony of others is reliable enough to be worthy of our actual trust?

Several philosophers have taken up this challenge. According to David Hume, to take a famous example, the basis of our trust in testimony is the observed 'constant conjunction' of testimony and fact. We rely on what others say since we have been able to observe that their reports are mostly correct. Thus, Hume's approach is in a sense reductive: rather than taking testimony to be an autonomous

source of knowledge on a par with perception, it regards the latter as more fundamental and as providing the rationale for the former. One could also regard testimony as valid a priori or simply as a fundamental source of knowledge that cannot be vindicated in terms of something else, just to mention two other possibilities that come to mind.

This chapter will be devoted to C. A. J. Coady's ambitious attempt to vindicate our testimonial reliance in his celebrated book from 1992.¹ The main target of the critical part of his essay is indeed Hume whose reductionism he rejects in favour of an attempted coherence justification of our trust in the word of others.

Coady's coherence argument is in several respects remarkably similar to Lewis's attempted validation of memory, what might well be a sheer coincidence considering the absence of any references to Lewis in Coady's book. Thus, the first task Coady sets himself is to establish a positive degree of credibility of testimony so as to ensure that someone's testifying to the truth of a given proposition makes it more likely to be true than it was before. This is supposedly accomplished by means of a transcendental argument reminiscent of one line of thought we were able to isolate in our discussion of Lewis. Some credibility of testimony is, we are told, a condition for the very possibility of testimony as an institution. Coady proceeds to argue, again in close parallel to Lewis, that considerations of coherence can, so to speak, amplify that already certified initial credibility to the degree that our reliance on the reports of others becomes warranted.²

Not surprisingly, some of the difficulties that we encountered in Lewis's theory also threaten Coady's enterprise. This holds, in particular, for the notorious problem of assigning a definite degree to the individual credibility (or we will not know how much coherence is necessary to raise the likelihood of truth to the level required for reliance), but Coady's discussion of coherence also raises some specific problems that lack counterparts in Lewis's vindication attempt.

¹ Unless otherwise indicated, page references in this chapter are to Coady (1992).

² As Fumerton (1995) notes, there is some uncertainty as to what Coady is really claiming in positive terms. At one point he seems to say that we can know a priori that testimony is generally reliable (96). He can also be read as suggesting that testimony is a fundamental source of evidence on a par with perception and memory (145). Nonetheless, the coherence argument is the only one that Coady develops in some detail, and the fact that this is carried out at the end of the part of his book called 'The Solution' strongly indicates that it represents his final word on the matter.

It is worth noting that the coherence argument, which will be central to this study, seems to have gone largely unnoticed in the already quite extensive secondary literature on Coady, which tends to focus instead on his criticism of Hume and his attempted 'Davidsonian' *reductio ad absurdum* of the notion that testimony could be entirely unreliable (see section 5.2. below).³ To the best of my knowledge, there is no detailed published discussion on Coady's final coherence argument, although the latter plays a pivotal role in his positive theory.⁴

In the following I will not engage in detailed criticism of the Davidsonian argument (for more on Davidson, see section 9.2). Rather my conclusion will be that Coady's overall argument fails even if the Davidsonian strategy is ultimately successful. Along the way I will attempt to establish the following: (a) the Davidsonian argument, if correct, provides too little information about the initial credibility of testimony to be useful as a basis for coherence reasoning; (b) Coady misconstrues the effect of what he calls 'cohesion', thinking that it would have a positive effect on our evaluation of reliability when in fact its effect is entirely negative; and (c) he falls prey to an ambiguity that is no less grave than the one he observes in Hume's treatment of the subject. Coady's point of departure, again, is his criticism of Hume's proposal that we are allowed to rely on the word of others so far as we have observed a reliable connection between testimony and fact testified, and this is also where I will begin.

5.2 Hume's 'Fatal Ambiguity'

Hume is of course well aware of the practical importance of being able to rely on others: In his own words, 'there is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men and the reports of eye-witnesses and spectators' (1978: 88). Given the significance of this practice, how can it be rationalized philosophically?

³ See for instance Graham (2000), Insole (2000), and Lyons (1997).

⁴ The coherence argument is addressed briefly in Fricker (1995: 410), where it is endorsed as a 'powerful line of argument'.

Hume offers an inductive justification that draws on his famous analysis of causality in terms of the ‘constant conjunction’ of cause and effect:

This species of reasoning, perhaps, one may deny to be founded on the relation of cause and effect. I shall not dispute about a word. It will be sufficient to observe that our assurance in any argument of this kind is derived from no other principle than our observation of the veracity of human testimony, and of the usual conformity of facts to the reports of witnesses. It being a general maxim, that no objects have any discoverable connexion together, and that all the inferences, which we can draw from one to another, are founded merely on our experience of their constant and regular conjunction; it is evident that we ought not to make an exception to this maxim in favour of human testimony, whose connexion with any event seems, in itself, as little necessary as any other. (1978: 111)

A few pages later, we find the following concise formulation: ‘The reason why we place any credit in witnesses and historians, is not derived from any *connexion*, which we perceive a priori, between testimony and reality, but because we are accustomed to find a conformity between them’ (1978: 113). In a nutshell, the reason why we place trust in testimony is, Hume thinks, because experience has shown it to be reliable.

Coady levels several objections at Hume’s contention, one centring on a ‘fatal ambiguity’ in the latter’s use of terms like ‘experience’ and ‘observation’ (80). These terms can either refer to the experience and observation of a single individual, in which case Hume’s thesis is ‘plainly false’, or they can refer to common experience and observation of mankind, in which case the contention is ‘question-begging’. Let us see why this is supposed to be so.

As for the first horn of the dilemma, suppose that we rely on others because each of us has observed for himself or herself a constant conjunction between what people report and the way the world is, so that each of us has good inductive grounds to expect the concurrence to continue into the future. Coady’s reply to this version of Hume’s thesis is, in its essence, that it seems to require too much of the capacity of single individuals to engage in extensive ‘field-work’ (82): ‘many of us have never seen a baby born, nor have most of us examined the circulation of the blood nor the actual geography of the world nor any fair sample of laws of the land, nor have we made the

observations that lie behind our knowledge that the lights in the sky are heavenly bodies immensely distant nor a vast number of other observations that [this reading of Hume's claim] would seem to require' (ibid.). The point is that we all rely on testimonies in areas where we have, as a matter of fact, not made any direct observations of our own and where we are *a fortiori* not in a position to confirm or disconfirm constant conjunctions of reports and facts reported. In Coady's own words, 'we rightly accept testimony without ever having engaged in the sort of checking of reports against personal observation that [Hume's thesis on the current interpretation] demands' (83).

As for the second horn, it is clearly fishy to support the concurrence of testimony and fact by citing as evidence 'our experience of their constant and regular conjunction' where 'our experience' refers to the common experience of mankind. Presumably we have access to this common experience only through our reliance on the testimony of others. On this reading, Hume is involved in vicious circularity since his attempt to provide a rationale for our reliance on others turns out to depend crucially on that very reliance.

5.3 The Argument for Individual Credibility

I will now proceed to what Coady takes to be a more fundamental trouble with the Humean account. In saying that we rely on testimony because we have found a constant conjunction between what it expresses and the way the world is, Hume presupposes that we can identify and understand testimonies independently of whether or not they are reliable. Outgoing from an initial state of ignorance as to the reliability, we are supposed to be able to identify a given utterance as a testimony whose correspondence with reality can then, in a second step, be confirmed or disconfirmed. Testimony is to be relied upon if a sufficient number of confirming (and no, or very few, disconfirming) observations have been made. As Coady correctly observes, this view implies the conceptual possibility of there being no observed conformity at all between testimony and reality. But, says Coady, it is highly doubtful whether this is indeed a possibility.

In an attempt at a *reductio ad absurdum*, he asks us to imagine a community of Martians that has the practice of reporting but where

there is no conformity between testimony and fact. The Martians are assumed, initially, to have a language that we can translate with names for distinguishable things in their environment and be in possession of 'suitable predicative equipments'. We find however, to our astonishment, that whenever they construct sentences addressed to each other in the absence (from their vicinity) of the things designated by the names, they seem to say what we (more fortunately placed) can observe to be false. But in such a situation the Martians would have excellent reasons *not* to rely on the 'reportive utterances' of others. The significance of this fact lies in one of its implications, namely that 'the Martian community cannot reasonably be held to have the practice of reporting' (87). For it is part of that practice that there be a tendency to rely on what has been reported, just as it is part of the practice of giving orders that those orders be regularly obeyed (*ibid.*). Martian 'reporting', as here described, would be a pointless game rather than a genuine speech act.

Not only would we have reasons to think that the Martians lack the practice of reporting. Even if we allow, for the sake of argument, that this practice exists in their community, we would not be in a position to find out empirically that there is a complete lack of correlation between what they say and the way the world is, unless we can understand what Martian reports actually say in the first place. But this is precisely where fundamental difficulties crop up, for given the circumstances as described it could even be questioned whether we really understand the contents of their utterances. Coady argues, not unconvincingly, that the situation imagined would lead to 'linguistic chaos' (89). In order to learn the meaning of Martian words and sentences we would have to have access to true Martian reports that, together with clues from the environment, would constitute the empirical basis for conjectures as to the meaning of Martian utterances. If, for instance, a Martian utters 'Kar do gnos u grin' while pointing its finger (or, if the Martian anatomy does not include fingers, some other suitable part of its constitution) to a tree in the garden we might take this as a basis for a preliminary guess that the utterance means 'There is a tree in the garden'. In the absence of reports that can be assumed correct we would have a hard time identifying the meaning of Martian words and sentences and the eventual conclusion might well be that the Martians are totally incomprehensible to us.

In the absence of a connection between reports and reality, we, the outsiders, would not be the only ones having severe difficulties figuring out what Martian utterances mean; the Martians themselves would be similarly disadvantaged and it would remain a mystery how their children could ever learn to master the Martian language. Much language instruction proceeds in the form of reports on the meaning of words, and so at least this part of the Martians' assumed practice of reporting would have to be in working order for the Martians to have anything like a public language that could be passed on to future generations. But, if so, how could it be that the hypothesized massive general breakdown in reporting does not affect reports on the proper use of language? The upshot is that under the imagined circumstances the very idea of a public language would be undermined. As Coady himself notices, his argument bears close resemblance to a line of reasoning made famous by Donald Davidson.

Let us now summarize Coady's Davidsonian argument. Given a general unreliability of 'reporting', (a) there could be no such things as reports, (b) even if there were reports, there could be no way of establishing Humean correlations or non-correlations since there could be no way of determining the contents of the alleged reports in order to correlate them, and so (c) the very idea of a public language would seem undermined.

Coady's argument, if correct, would indeed establish the impossibility of a complete and utter unreliability of testimony, but it remains somewhat unclear what positive conclusions could be derived from that reasoning. Coady's cautious conclusion is that 'there must be at least the *minimum connection* between testimony and reality that the breakdown of the no-correlation possibility reveals' (96, my italics). There is at least some, however small, positive correlation between testimony and reality but, for all the argument shows, this connection need not be very pronounced. But he also states, more boldly, that '[t]he above discussion itself strongly suggests a *reasonable degree of reliability* about the testimony of others, whether they be aliens or natives' (168, my italics). And he even goes as far as saying: 'From what our discussion of that breakdown exhibited we may well conclude that the connection has to be *quite extensive*' (my italics), the reason being that '[i]f, as I claimed earlier, the ability to use language meaningfully is connected with the

making of true reports then it is surely the *consistent* making of true reports that matters'. By the same token, he claims to have shown that 'an *extensive commitment* to trusting the reports of others was a precondition for understanding their speech at all' (176, my italics).

Interestingly, when it really matters in his subsequent argumentation Coady relies—wisely, I believe—only upon the weak conclusion that testimonies must be credible at least to some degree, however small that degree might be. He notes, in particular, that the argument against Hume 'is effective if it demonstrates merely that linguistic communication commits a person to *some* degree of trust in the word of others' (153). He continues: 'This would, however, be too limited a conclusion if the argument is to serve the positive role sketched above [i.e. the role of vindicating our actual trust].' Coady is here saying, or implying, that the argument against Hume can only support a minimal conclusion and that, therefore, a complete legitimization of our reliance on the word of others would require an extended argument. This puts Coady in exactly the same position as Lewis, who, we recall, had to concede that his argument for the positive credibility of memory did not indicate any particular degree thereof.

5.4 The Invocation of Cohesion and Coherence

Coady is not satisfied with having established a mere positive individual credibility of natural testimony since 'this leaves us with the question whether any sort of argument can be mounted to provide some justification or philosophical rationale for what is in fact our very extensive trust in testimony' (152). And '[i]n the absence of such an argument the thought may very naturally arise that, although we must trust some testimony, neither the extent nor centrality of our actual reliance is rationally supportable' (152–3). A constraint on such an argument is that it avoids 'succumbing to the individualist temptations that seem to have bedevilled other attempts at justification' (174–5). Rather, '[a]ny such project must begin by assuming the existence of a public language in which the testimony to be scrutinized is to be made available' (153). In response to this challenge, Coady suggests 'certain broad facts of cohesion and coherence imply the legitimacy of the strong commitment to trusting the

word of others that is embodied in our actual cognitive procedures' (176). Let us see how he arrives at this conclusion.

Coady begins by distinguishing different 'informational routes' (169). These are constituted by an individual's perceptions, memories, inferences, and 'learnings', the latter being his semi-technical term for information arrived at via the testimony of others. Coady now argues that the informational routes cannot be treated in 'an isolated atomistic fashion' (168), a remark that applies 'both to the way we get into a position to exercise the powers in question and to the assessment of their results' (168–9). As an example of the first phenomenon we may take an individual's present perception of something in front of him as an eighteenth-century mahogany architect's desk which 'will be determined not only by the gross perception of certain colours and shapes but by the memories and inferences of both himself and others which are built into the conceptual and perceptual skills with which he approaches this particular cognitive encounter' (169). Thus the outcome of the perceptual route will in part depend on the previous outcome of other routes, in this case memory and inference. Coady refers to this 'integration of informational routes' as 'cohesion', reserving the term 'coherence' for agreement between outcomes, as when what we are told concurs with what we remember.⁵

As an interesting effect of cohesion, it is not possible to reach a neutral base with respect to testimony without also giving up a lot of perceptions, memories, and inferences: 'Someone who sought to isolate an individualist basis in perception, memory and inference, in order to test the reliability of testimony, would not only face the problems of language and understanding discussed earlier in this chapter but would have to discount an enormous amount of what goes into normal perception, inference, and memory' (170).

One may wonder what the relationship might be between cohesion and coherence. In Coady's view, these concepts, though

⁵ Coady provides the following additional examples to illustrate how 'all our basic sources of information are closely interwoven' (1992: 99): 'we often correct memory by memory and correct or reject apparent observations which conflict with memories. So, I may seem to glimpse a friend in the corridor of a building but reject the evidence of sense because I recall that he is in another country. Here we have a more or less conscious intellectual process but the influence of memory on perception can be more direct, as when I believe that my friend is in another country and so actually fail to recognize him visually when, upon his unexpected return, he is before my eyes' (*ibid.*).

distinguishable, are closely related (169), but he remains vague as to the exact nature of the relationship, saying only obscurely that ‘our attitude to outcomes is all of a piece with that existing integration which makes for the operation of a particular channel of information’ (ibid.).

Having clarified the nature of cohesion and coherence, including our tolerance for some incoherence to be resolved in subsequent enquiry, Coady makes the following statement of their place in the epistemology of testimony:

None the less, coherence remains the ideal and our fairly generous capacity to tolerate incoherence or lack of fit testifies to the generally very satisfactory state of coherence which we continue to find and cheerfully expect between the different informational sources. The cohesion is not disturbed but reinforced by the way things continue to turn out. It is not that we somehow ‘prove’ from purely individual resources that testimony is generally reliable but that, beginning with an inevitable commitment to some degree of its reliability, we find this commitment strongly enforced and supported by the facts of cohesion and coherence. (173)

There are a number of critical remarks to be made at this point. First of all, an argument to the effect that our commitment to the word of others is ‘strongly enforced and supported’ by the supposed facts of cohesion and coherence falls short of establishing the stronger conclusion that testimony can be relied upon. In order to support that conclusion, it must be shown not merely that our commitment is *strengthened* by the observation of cohesion and coherence but that it is thereby made *strong enough* for the purposes of reliance.

Does Coady’s argument in fact support the stronger conclusion that we should assign a high enough degree of reliability to testimonies upon taking cohesion and coherence into account? I think not. Let us begin with coherence. As we saw in connection with Lewis, the eventual probability of reliability arrived at via Bayes’s theorem is contingent not merely on facts of coherence but also on the initial degree of reliability or credibility. If the latter cannot be fixed within reasonable limits, coherence alone need not be a very significant fact.

This is a point which Coady seems willing to endorse. In his discussion of single miracle testimonies, he notes that, as a consequence of Bayes’s theorem, ‘[i]f we view testimony as a very weak evidential reed and our fellow observers as by nature grossly

credulous we will come to very different conclusions from those whose fundamental epistemic outlook is more trusting' (192 n. 20). He goes on to say that the same observation applies to situations involving several testimonies in agreement where, as he writes, our 'attitude to reliability can play a significant role'.⁶

All that Coady has shown, before cohesion and coherence were invoked, is that there is *some* initial degree of credibility pertaining to testimonies as such, while conceding, at least in his more cautious moments, that we cannot say with confidence what that degree is. So, what follows from Coady's reasoning is at best, first, that we have reason to assign a positive degree of credibility to testimony before taking coherence into account; second, that the degree of credibility is raised by taking coherence into account; but it does not show that we are in a position to say how high it gets and, in particular, whether it is high enough to warrant our actual full and unreserved trust in normal cases. To know this, we would have to know, at least within reasonable limits, what the initial degree of credibility was, but on this crucial point Coady's argument is silent.

Turning to cohesion, one may question whether it contributes at all to our confidence in the reliability of testimony. The presence of cohesion between informational routes actually means that these routes are dependent on each other in the evidential sense. If what we perceive is dependent on what we remember and infer, as Coady's own example with the mahogany desk indicates, then it will not come as a great surprise that our perceptions agree with our memories and inferences. Facts of cohesion actually speak against, and not in favour of, any enforced reliability commitment as the effect of observing agreement between outcomes. In the extreme case where one route repeats the outcome of another, agreement has no effect at all on the probability of reliability. It is curious that Coady mentions cohesion as a factor that would support the reliability of testimony in the face of agreement between outcomes. If my acceptance of what others tell me depends on my own

⁶ '[T]here is an argument that one's fundamental attitude to testimony may be unimportant in some particular cases. Corroboration is the most typical sort of case. I may regard each of ten witnesses as highly unreliable but accept their joint testimony on the grounds that its truth best explains their agreement. Yet even here, I should say, attitude to reliability can play a significant role, since it may favour an alternative explanation of the convergence, such as collusion. The "sceptic" will be more inclined to seek out and countenance such alternatives' (191 n. 20).

perception and memories, agreement will not be as significant as it would have been, had there not been this sort of dependence. Similarly, if what I (think I) hear depends on what I (think I) see. It is an empirical question how extensive this cohesion is between informational routes. The main point here is just that, unlike what Coady seems to think, a high degree of such cohesion would speak against and not in favour of his justificatory enterprise.⁷

These are severe shortcomings, especially the problem of assigning an initial degree of credibility. And yet there is more to come by way of criticism. In the next section, I will argue that Coady, in the passage quoted on the supposed virtues of cohesion and coherence, relies on an ambiguity similar to the one he accused Hume of succumbing to and hence faces a similar Humean dilemma.

5.5 Coady's Fatal Ambiguity

Let us return to the ambiguity inherent in Hume's argument for the reliability of testimony, which referred to an allegedly observed constant conjunction between report and fact. As Coady pointed out, there are two readings of that argument, depending on whether by 'observation' is meant the common experience of humanity or the mere solitary observation of David Hume. On the first reading, Hume's argument is question-begging because we would have to rely on the testimonies of others in order to find out what they have experienced. On the second reading, the argument seems plainly wrong, since none of us, Hume himself included, has done the amount of 'field-work' that would be necessary to vindicate trust in testimony. We trust testimonies in areas where we have little or no chance of making observations of our own. Many of us have not directly witnessed 'a baby born', 'the circulation of the blood', 'the actual geography of the world', nor have we been in a position to verify for ourselves 'that the lights in the sky are heavenly bodies immensely distant', and so on (82).

But compare Hume's discussion of testimony and fact and his suspicious-looking reference to 'our experience of their constant

⁷ I return to Coady's concept of cohesion in section 10.2, then in connection with Jamesian wagering as an answer to scepticism.

and regular conjunction' with the following already quoted passage from Coady:

None the less, coherence remains the ideal and our fairly generous capacity to tolerate incoherence or lack of fit testifies to the generally very satisfactory state of coherence which *we continue to find and cheerfully expect between the different informational sources*. The cohesion is not disturbed but reinforced by the way things continue to turn out. It is not that *we somehow 'prove'* from purely individual resources that testimony is generally reliable but that, beginning with an inevitable commitment to some degree of its reliability, *we find this commitment strongly enforced and supported by the facts of cohesion and coherence*. (173, my italics)

Is this not ambiguous in just the same way as Hume's statement was seen to be? Take, for instance, the claim that 'we find this commitment strongly enforced and supported by the facts of cohesion and coherence'. This can be read as referring either to humanity's common experience of facts of cohesion and coherence or, alternatively, to the solitary coherence experience of Coady. When read in the first way, it is clearly question-begging, for how could we find out about others' experiences of coherence without asking them and relying on their answers? Assuming that we could do this is not allowed in this context, since Coady's argument is intended to provide us with a vindication of our reliance on testimony and must not presuppose that such trust has already been secured.

Hence, we must turn to the other reading in terms of individual experiences of coherence. Is it correct that each individual has observed the coherence of testimonies of others with the outcome of his or her other 'informational routes', like perception, memory, and inference? Let us begin with perception. Take a person who has not been able to verify perceptually statements about the circulation of the blood, the actual geography of the world, astronomical facts, etc., but who nonetheless relies on testimony in these areas. It seems clear that there are such persons, and this was also a premiss upon which Coady's argument against Hume relied. Now such a person clearly would not be able to check the coherence of testimonies in this area with any outcome of the perceptual route since there is no such outcome. But perhaps she could check the validity of such testimonies against her memories and inferences. But whatever the source of those memories and inferences might be, they cannot

derive from our first-hand perceptions of such facts for, by assumption, we have no such perceptions. Hence, these memories and inferences must be mere fabrications, without any contact whatsoever with reality. But if so, they will be entirely unreliable, when taken singly, and for reasons that should be familiar by now their coherence will contribute nothing to the eventual probability of reliability.

In summary, we have seen that Coady's attempted final validation of testimony is strikingly similar to Lewis's justification of memory. In particular, they both leave unresolved the notorious problem of how to assign individual credibility. Furthermore, Coady seriously misrepresents the effect of cohesion, which is essentially causing evidential dependence. Finally, Coady's own constructive proposal falls prey to exactly the same type of ambiguity that we found in Hume's treatment of the subject.

5.6 Closing Remarks on the Anti-sceptical Use of Coherence

It is time for a short summary of Part I. What our discussion has revealed, I hope, is that two of the most celebrated and detailed coherence-based theories—Lewis's and Bonjour's—are simply probabilistically unsound. Bonjour's coherence theory is founded on the false probabilistic assumption that individually useless data can, if they are independent, become collectively useful, in the sense of yielding, when combined, a high posterior probability. Lewis wisely rejected this notion, pointing out that coherence has an effect only if independent data have some degree of individual credibility.

But Lewis, too, was unrealistically optimistic in his assessment of what coherence can actually accomplish. His weak foundationalism—his attempt to provide a 'final validation of empirical knowledge by reference to ultimate data in some sense presently given, and to the congruence of such data' (356), given our supposed partial ignorance regarding the initial credibility of those data—shares with Bonjour's theory the deficiency of being founded on an incorrect probabilistic supposition. In Lewis's case, the erroneous, but essential, assumption is that posterior probability is determined by coherence alone once the minimum conditions of individual credibility and independence

are satisfied. Thus, Bonjour's coherentism and Lewis's weak foundationalism turn out to be untenable already on the grounds that they rely, each in its own essential way, on probabilistic assumptions that are simply incorrect. As we just saw, our central criticism of Lewis carries over unmodified to C. A. J. Coady's attempted vindication of natural testimony.

There is a further probabilistic shortcoming that Lewis and Bonjour both succumb to. They both presuppose, falsely, that the issue of individual credibility can be discussed separately from the issue of which hypotheses are possible as regards the reliability of reports. As we have seen, by contrast, once the set of possible hypotheses has been specified, the individual credibility of reports is determined by the laws of probability.

To see the point, suppose that X_1, \dots, X_n are mutually exclusive and exhaustive hypotheses regarding the reliability. They can be, for example, complete reliability, randomization, and various forms of lying. Given this specification of alternatives, the individual credibility of a report E with content H is given by

$$P(H/E) = P(H/X_1, E)P(X_1/E) + \dots + P(H/X_n, E)P(X_n/E)$$

It is often plausible to assume that the different reliability hypotheses are independent of the report E , in which case we get:⁸

$$P(H/E) = P(H/X_1, E)P(X_1) + \dots + P(H/X_n, E)P(X_n)$$

In Lewis's case, the failure to appreciate this relationship between reliability hypotheses and individual credibility proved harmless. The reason is that plugging the hypotheses that are in his view serious alternatives—reliability and randomness—into the equation above has the effect of making $P(H/E)$ exceed $P(H)$, just as Lewis insisted should be the case.

For Bonjour's, by contrast, the shortcoming has serious consequences, as it creates an incompatibility problem. The problem is caused by the fact that individual credibility is ensured (given some eminently plausible assumption), even if we add, as Bonjour urges that we should, 'lying' to the list of possible hypotheses alongside

⁸ This may not be so plausible if the report E is initially extremely unlikely. If someone presents me with a very implausible report, I tend to revise my judgement that the person is reliable. This complication does not affect the point made here.

reliability and randomness. This is in blatant conflict with what Bonjour sees as the hallmark of his coherentism, i.e. his insistence that cognitively spontaneous beliefs need have no individual credibility. We noted also another shortcoming in Bonjour's account: the clash between his explicit statement that cognitively spontaneous beliefs can be independent and his admission of collusion as a serious initial possibility.

The coherence theory is impotent as a reply to radical scepticism. But that may not be a catastrophic fact. I will argue, in Chapter 10, that we have no reason to take the sceptic's challenge seriously in any case. The aim of this essay, however, is not primarily to penetrate the problem of radical scepticism, but to enquire into the relation, if there is any, between coherence and truth. In Part II, I will turn to the comparative claim that more coherence implies higher likelihood of truth.

Part II

**Does More Coherence Imply Higher
Likelihood of Truth?**

This page intentionally left blank

Making the Question Precise

6.1 Degrees of Coherence

Coherence does not by itself imply high likelihood of truth. The anti-sceptical use of coherence is based on a false assumption. But if we cannot show that coherence implies high likelihood of truth, can we at least make it plausible that *more* coherence implies *higher* likelihood of truth?¹ This question was posed in a seminal paper by Peter Klein and Ted A. Warfield, which triggered a lively debate in the journal *Analysis*.² The question then is this: if a system *S* is more coherent than another system *S'*, are we then allowed to conclude that *S* is more likely than *S'* to be true as a whole? If the answer is in the affirmative, we will say, following Klein and Warfield, that coherence is *truth conducive*.

How important, philosophically, is the comparative question? BonJour, for one, takes it to be central to the anti-sceptical project:

Finally, coherence is obviously, on any reasonable view, a matter of degree (as is stability). Hence the conclusion of the envisaged argument [for BonJour's coherence theory] should be that *the likelihood that a system of beliefs corresponds to reality varies in proportion to its degree of coherence* (and stability) other things being equal. (BonJour 1985: 170, my italics)

BonJour is here saying that one urgent project for the coherence theorist is to show that a higher degree of coherence and stability implies a higher likelihood of truth. I will follow Klein and Warfield and most other authors who have addressed the comparative question

¹ Actually, this question is slightly misleading, though in an innocent way, since it suggests that the comparative statement is weaker than the absolute statement—that 'coherence implies a high likelihood of truth' implies but is not implied by 'more coherence implies a higher likelihood of truth'. In reality, the two statements are logically independent.

² Klein and Warfield (1994; 1996), Merricks (1995), Cross (1999), Shogenji (1999; 2001), Akiba (2000), and Olsson (2001).

in disregarding the stability issue.³ This is due to the fact that we share their aim to ascertain whether coherence *by itself* is in some interesting sense connected with likelihood of truth. While the truth conduciveness claim that is considered here does not match exactly that made by BonJour, the relevance of our study to his theory should be clear. If more coherence does not have the said effect in isolation—as I will indeed argue later on—there is little reason to believe that it would have it in the presence of stability. At the very least, the burden of proof would be on BonJour to make it plausible that coherence plus stability has the benefits which coherence alone lacks.⁴

Already in Chapter 2 we encountered the problem of assigning a precise meaning to the concept of coherence. This led to the suggestion that agreement is a species of coherence. However, this is not of much help in the present context which focuses on comparative uses of coherence. It is true that BonJour made some tentative proposals in this direction in his 1985 book, but we also saw that they leave much to be desired, and in a more recent publication BonJour concedes that ‘the precise nature of coherence remains an unsolved problem’ (1999: 24). The reason is that

spelling out the details of this idea, particularly in a way that would allow reasonably precise assessments of *comparative coherence*, is extremely difficult, at least partly because such an account will depend on the correct account of a number of more specific and still inadequately understood topics, such as induction, confirmation, probability, explanation and various issues in logic . . . (ibid.: 123, my italics)

BonJour, more than a decade after his main book appeared, takes the comparative notion of coherence to be central, albeit ‘extremely difficult’ to define. In the present part of this essay, I will go one step further than BonJour and argue that the task of spelling out the details of comparative coherence is not only extremely difficult; there are strong indications that it is downright impossible.

This chapter will be devoted to making the comparative question precise. This involves clarifying the two crucial notions: that of a degree of coherence and that of truth conduciveness. Those concepts

³ For an exception, see Cross (1999). ⁴ Cf. section 4.1.

will be explicated in a way that is in line with the Lewis–BonJour tradition. What I hope to arrive at is not only a historically adequate but also a charitable rendering of the coherence theory. As a preliminary I will argue that, once the coherence theory is understood along the lines drawn here, several arguments that have been levelled in the literature against the truth conduciveness of coherence lose their force.

We recall our explication of Lewis’s ‘set of supposed facts asserted’ in terms of a testimonial system $\mathbf{S} = \{\langle E_1, A_1 \rangle, \dots, \langle E_n, A_n \rangle\}$ where E_i is a report to the effect that A_i is true. We will say that A_i is the *contents* of report E_i . The *contents* of a testimonial system $\mathbf{S} = \{\langle E_1, A_1 \rangle, \dots, \langle E_n, A_n \rangle\}$ is the ordered set of report contents $\langle A_1, \dots, A_n \rangle$. By the *degree of coherence* $C(\mathbf{S})$ of such a testimonial system we will mean the degree of coherence of *its content*. This reflects the pivotal idea that coherence is a property on the content level, a principle which will play a crucial role in the impossibility theorem presented at the end of Chapter 7 and proved in Appendix B.

In this section, the strategy will be to explore what a coherence measure could look like as defined for sequences of propositions. Once we have a measure of coherence defined over sequences of propositions, we can define the degree of coherence of a testimonial system as the degree of coherence of the sequence of its content propositions.

One structural requirement of a coherence measure should be mentioned at the outset, namely its independence of the order in which the elements of the given set are listed or, in more technical language, its invariance under arbitrary permutations of the supposed facts asserted. Thus, we require of any coherence measure C worthy of the name that $C(\langle A_1, A_2, \dots, A_n \rangle) = C(\langle B_1, B_2, \dots, B_n \rangle)$ whenever $\langle B_1, B_2, \dots, B_n \rangle$ is a permutation of $\langle A_1, A_2, \dots, A_n \rangle$. All measures that will be proposed below satisfy that condition.

We recall Lewis’s definition of congruence in terms of supposed facts ‘so related that the antecedent probability of any one of them will be increased if the remainder of the set can be assumed as given premises’ and also that this definition is inadequate in general. Congruence in Lewis’s sense does not correspond to our intuitive notion of coherence in cases of more than two testimonies. This was shown in relation to the octogenarian example. It will prove useful to

recall the details. We assumed there to be a reasonable number of students and a reasonable number of octogenarians (80–89-year-olds) and, further, that all and only students like to party and that all and only octogenarians are birdwatchers, and, finally, that there are some, but very few, octogenarian students. We then called attention to the set consisting of $A_1 =$ ‘The suspect is a student’, $A_2 =$ ‘The suspect likes to party’, $A_3 =$ ‘The suspect is an octogenarian’, and $A_4 =$ ‘The suspect likes to watch birds’, noting that while this set is congruent in Lewis’s sense, it cannot be said to be coherent in a pre-systematic sense: one half of the story (the one about the partying student) is very unlikely given the other half (the one about the birdwatching octogenarian).

What has gone wrong here? One notable thing is that the joint probability of the propositions in the set is very low; the probability that the suspect is both a partying student and an octogenarian birdwatcher is close to zero. This suggests identifying the degree of coherence of a set with its joint probability:

$$C_0(A, B) = P(A \wedge B)$$

This measure takes on a minimum value of 0 if and only if there is no overlap between A and B . It takes on a maximum value of 1 just in case $P(A) = P(B) = 1$, i.e. just in case both A and B are for sure. It is obvious how to generalize this measure to the case of an arbitrary (finite) number of propositions:⁵

$$C_0(A_1, \dots, A_n) = P(A_1 \wedge \dots \wedge A_n)$$

This, however, is not the only plausible way to measure coherence. Suppose that there has been a robbery. To get an unbiased view on who might have committed the crime you decide to consult four different witnesses. Suppose that the first two witnesses both claim that Steve did it, the third witness that Steve, Martin, or David did it, and the fourth witness that Steve, John, or James did it. Which pair of statements is the more coherent—that delivered by the first two sources or that delivered by the last two sources? It is difficult to escape the feeling that the statements delivered by the first two sources are more coherent in an intuitive sense. After all, unlike the

⁵ I will be concerned throughout exclusively with the coherence and truth of *finite* sets.

last two witnesses the first two say exactly the same thing, so that the degree of agreement is higher. According to the C_0 -measure, by contrast, the degree of coherence is the same in the two cases, equaling the probability that Steve did it.

This leads us directly to the next proposal. Rather than measuring the overlap of propositions, we might measure the extent to which they agree. The more the propositions agree, the more coherent they are. From this perspective, propositions that coincide are always maximally coherent, even if their joint probability is not very high. A simple way to measure the extent of agreement is the following (Olsson 2002a, Glass 2002):

$$C_1(A,B) = \frac{P(A \wedge B)}{P(A \vee B)}$$

$C_1(A,B)$ measures how much of the total probability mass assigned to either A or B falls into their intersection. $C_1(A,B)$ takes on values between 0 and 1. As before, the degree of coherence is 0 if and only if $P(A \wedge B) = 0$, i.e. just in case A and B do not overlap at all, while the degree of coherence equals 1 if and only if $P(A \wedge B) = P(A \vee B)$, i.e. just in case A and B coincide. The measure is straightforwardly generalizable:

$$C_1(A_1, \dots, A_n) = \frac{P(A_1 \wedge \dots \wedge A_n)}{P(A_1 \vee \dots \vee A_n)}$$

A potential drawback of this measure is its assignment of the same coherence value, namely 1, to all cases of total agreement, regardless of the number of witnesses that are involved. Against this, it may be objected that agreement among the many is more coherent than agreement among the few.

The following alternative coherence measure was introduced by Tomoji Shogenji (1999):

$$C_2(A,B) = \frac{P(A/B)}{P(A)} = \frac{P(A \wedge B)}{P(A) \times P(B)}$$

It is easy to see that this measure is sensitive to the number of reports in cases of total agreement: n agreeing reports correspond to a coherence value of $1/P(A)^{n-1}$. Thus, if coherence is measured in this fashion, then agreement among the many comes out as more

coherent than agreement among the few. Like the other measures, $C_2(A,B)$ equals 0 if and only if A and B do not overlap. Shogenji has proposed the following generalization:

$$C_2(A_1, \dots, A_n) = \frac{P(A_1 \wedge \dots \wedge A_n)}{P(A_1) \times \dots \times P(A_n)}$$

Another interesting difference between C_1 and C_2 concerns their sensitivity to the specificity of the asserted facts or, what comes to the same thing, their prior probability. Let us modify the robbery example somewhat. Suppose, as before, that the first two witnesses report that Steve did it, and that the third witness reports that Steve, Martin, or David did it. But suppose now that the report of the fourth witness coincides with that of the third, so that they both report that Steve, Martin, or David did it. Which testimonies are now most coherent, those of the first two witnesses or those of the last two? Both pairs are in full agreement, and so their degree of coherence should presumably be the same. Appealing to our intuitions of mutual support yields the same result; for each set, the one statement in the set is established if the other is assumed as given premiss. This is also what we get if we apply C_1 .⁶ On the other hand, the agreement between the first two testimonies is surely much more striking since, unlike the last two, they coincide on a very specific statement. Therefore, one could argue, the degree of coherence should also be greater in the first case. This is also what C_2 yields.

There may be many other ways of measuring the degree of coherence or agreement. While both C_1 and C_2 have some initial appeal, I doubt that our intuitions are clear enough to single them out as the only plausible candidates. Nonetheless, they do provide us with a useful starting point for concrete discussion.

In general, we will mean by a (probabilistic) *coherence measure* as defined for sequences of propositions any numerical measure $C(A_1, \dots, A_n)$ defined solely in terms of the probability of A_1, \dots, A_n (and their Boolean combinations) and standard arithmetical operations. The measures introduced so far are, of course, all cases in point.

⁶ There is textual evidence suggesting that C. I. Lewis construed congruence as not varying with detail or specificity: 'The coherence of a novel, or of the daydreams we are aware of fabricating as we go along, can never have the slightest weight toward crediting the content of them as fact, no matter how detailed and mutually congruent such items may be' (1946: 357, my emphasis).

When it matters, we will make explicit the dependence of a measure of coherence on a probability distribution, writing $C_P(A_1, \dots, A_n)$.

6.2 Coherence and Logical Closure

In his contribution to the *Analysis* debate, Ken Akiba (2000) levels a number of objections at the very idea of defining coherence probabilistically. His specific target is Shogenji's measure of coherence, the one we have labelled C_2 . In this section, I will explain why what I take to be his most serious objection, while it may apply to other conceptions of coherence, is inconsequential once coherence is construed as a property of testimonial systems.

Akiba reasons around a simple example in which we are assumed to be wondering what number a die will show the next time it is cast. Consider the following propositions:

$B(2)$ = The die will show 2.

$B(2,4)$ = The die will show 2 or 4.

$B(2,4,6)$ = The die will show 2, 4, or 6.

Akiba notes that $C_2(B(2), B(2,4)) = 3$ and $C_2(B(2), B(2,4,6)) = 2$. He goes on to make the following remark:

[W]e want to say that the coherence of $\{B(2), B(2,4)\}$ and $\{B(2), B(2,4,6)\}$ should be no different from the (self-)coherence of $B(2)$ (or the singleton $\{B(2)\}$), for $B(2,4)$ and $B(2,4,6)$ are both just logical consequences of $B(2)$, so whoever believes $B(2)$ may as well believe $\{B(2), B(2,4)\}$ and $\{B(2), B(2,4,6)\}$. However, if you apply $[C_2]$ to the case in which N [the number of propositions] = 1, $C_2(B(2)) = P(B(2))/P(B(2)) = 1$, not 2 or 3. (ibid.: 357, notation adapted)

Disregarding the troublesome application of coherence to singletons (see Chapter 2 for a discussion), Akiba is here making essentially two claims. First, he is suggesting that a person believes, or should believe, in the logical consequences of his or her beliefs. I have no quarrel with this. His second point is that sets that have the same logical consequences should be assigned the same degree of coherence. Thus, $\{B(2), B(2,4)\}$ and $\{B(2), B(2,4,6)\}$ should come out as equally coherent. Since Shogenji's measure gives a different result, it is, in Akiba's view, inadequate.

But why should we assign the same degree of coherence to sets that have the same logical contents? Perhaps Akiba is suggesting that we view such sets as being identical. So, for instance, $\{B(2), B(2,4)\}$ and $\{B(2), B(2,4,6)\}$ should not be treated as two different belief systems but as one and the same belief system. Because they are identical, they should be assigned the same degree of coherence.

Akiba's proposal amounts, in effect, to equating a belief system with its logical closure, i.e. the system obtained by adding to the original set all its logical consequences. This is a common way to represent beliefs in formal epistemology, especially in the literature on belief revision (e.g. Levi 1980). Thus understood, Akiba's point is that the idea of coherence collapses in the context of a closed set representation of belief systems.

With this I agree. Sven Ove Hansson and I made essentially the same point in Hansson and Olsson (1998). Unlike Akiba, however, we did not conclude from this collapse that it does not make sense to define coherence. Instead, we saw it as raising the question of whether the concept of coherence can be made intelligible for non-closed sets.

In fact, as I have argued in Part I of this book, there is ample evidence in the work of Lewis, Bonjour, and other coherence theorists that they are concerned not with all beliefs—derived or non-derived—but only with those that are non-derived. We recall Lewis's 'sets of supposed facts asserted' and Bonjour's cognitively spontaneous beliefs. The totality of beliefs that have an independent standing in this sense, and are not merely derived from other beliefs, will not form a logically closed set.

The upshot is that, unlike what Akiba thinks, it is uncharitable to dismiss the coherence theory because of difficulties that arise when coherence is applied to logically closed sets. To the best of my knowledge, no important coherence theorist has seriously considered this use of their central concept.

6.3 Testimonial Truth Conduciveness

Our question is whether coherence is truth conducive in the sense that more coherence implies a higher likelihood of truth. We have already said that coherence applies first and foremost to testimonial systems.

Thus the question we need to answer is this: if testimonial system $\mathbf{S} = \{\langle E_1, A_1 \rangle, \dots, \langle E_n, A_n \rangle\}$ is more coherent than testimonial system $\mathbf{S}' = \{\langle F_1, B_1 \rangle, \dots, \langle F_m, B_m \rangle\}$, does that mean that \mathbf{S} is also more likely than \mathbf{S}' to be true as a whole? By the *likelihood of truth* (probability) of a testimonial system $\mathbf{S} = \{\langle E_1, A_1 \rangle, \dots, \langle E_n, A_n \rangle\}$ we shall mean $P(\mathbf{S}) = P(A_1, \dots, A_n / E_1, \dots, E_n)$, i.e. $P(\mathbf{S})$ equals the joint probability of the contents of the testimonial system conditional on the testimonies.

Note that we cannot say that the likelihood of truth is simply $P(A_1, \dots, A_n)$ on pain of violating the requirement of total evidence, which says that all available evidence must be taken into consideration when computing probabilities. Since, by hypothesis, E_1, E_2, \dots, E_n constitute evidence (in the form of testimonies) for A_1, A_2, \dots, A_n , respectively, we must condition on the former when computing the likelihood of truth of the latter. We recall here Bonjour's Doxastic Presumption which assured that facts of belief can and should be taken as evidence for the purposes of epistemological enquiry.

At this point it is necessary to say something about the interpretation of probability. The only problematic cases are those involving *doxastic* systems, where the evidence is of the type 'S believes that A'. If one thinks, as I tend to do, that S is committed to assigning a probability of 1 to everything she fully believes, then one must also grant that the likelihood of truth or probability of S's doxastic system is always 1, if the probabilities are S's own relative to her current doxastic position. This trivialization can be avoided if S considers her beliefs from some weaker doxastic position, e.g. one at which no empirical propositions are assumed true. Another way to avoid it is to consider the subjective probabilities of some other person, and not those belonging to the subject herself. Finally, an objective interpretation of probability is possible. This is a genuine problem for an anti-sceptical coherence theory that focuses on doxastic systems, Bonjour's 1985 theory being the most prominent case in point.⁷

These considerations suggest that we define (testimonial) truth conduciveness as follows:

Definition: A coherence measure C is *truth conducive* if and only if: if $C_P(\mathbf{S}) > C_{P'}(\mathbf{S}')$, then $P(\mathbf{S}) > P'(\mathbf{S}')$.

⁷ For more on the interpretation of probability in this context, see Olsson (2002a).

Hence, a coherence measure is truth conducive whenever more coherence means higher likelihood, regardless of how probabilities are assigned and regardless of what systems are compared.

Why do we allow both the probability distribution and the testimonial system to vary between situations that are compared with respect to their relative degree of coherence? Well, why not? I am not aware of any reasons to keep the probability assessments fixed while varying only the testimonial system. By the same token, there seems to be no argument for fixing the testimonial systems while varying the probabilistic assumptions.

We can now pose our question about the truth conduciveness of coherence in a precise manner. What we want to know is whether there are any testimonially truth conducive measures of coherence. Answering this question turns out to be a rather complex matter that must be deferred to a separate chapter (Chapter 7). There I will also consider a weaker *ceteris paribus* conception of truth conduciveness. The remainder of this chapter will be devoted to an application of our account of truth conduciveness to a problem posed by Peter Klein and Ted Warfield.

Before entering that discussion I would like to add a few words on the relation between my account of truth conduciveness and one proposed by Charles B. Cross (1999). Cross makes his proposal in the context of a discussion of BonJour's theory. His account is less general than mine though, in that it is concerned exclusively with the doxastic case, where 'testimonies' come in the form of reports to the effect that a person believes this or that. Following Cross, we let the relation $\mathcal{J}(B, c, s, r)$ stand for 'B is justified to a degree defined by $\langle c, s, r \rangle$ '. The relation holds if and only if 'B is the conjunction of the members of the current belief set of an actual agent whose belief history has length r and consists of belief sets that have remained coherent to degree c and stable to degree s while satisfying the Observation Requirement' (ibid.: 189). In order to satisfy BonJour's Observation Requirement, a system of beliefs, as we saw in Part I, 'must contain laws attributing a high degree of reliability to a reasonable variety of cognitively spontaneous beliefs' (BonJour 1985: 144).

Cross goes on to construe BonJour's truth conduciveness claim as follows: If $\langle c_2, s_2, r_2 \rangle$ represents a greater degree of justification

than $\langle c_1, s_1, r_1 \rangle$, then $P(B_2/\mathcal{J}(B_2, c_2, s_2, r_2)) > P(B_1/\mathcal{J}(B_1, c_1, s_1, r_1))$. Note that Cross, too, conditions on the propositions being believed by an actual person. But this is not the only thing he conditions on. In fact, he explicates the truth conduciveness of a certain notion of (degree of) justification seen as applicable to a package of three parameters, including not only the degree of coherence but also two additional parameters: the degree of stability and the length of the belief history. While Cross may well be right in that this conception corresponds to Bonjour's intentions, he has not answered our question whether, and in what sense, coherence *per se* is truth conducive. Perhaps he did not intend to do so. At any rate, it seems to me sound methodology to study the truth conduciveness of the components before studying that of a whole package. The Bonjour–Cross strategy thus reverses the natural order of enquiry.

6.4 Why the Klein–Warfield Argument Fails

In their 1994 article, Klein and Warfield (K&W) argue that coherence is not truth conducive: more coherence does not imply higher likelihood of truth.⁸ They define an extension B' of a belief system B to be *non-trivial* just in case some of the beliefs that are in B' but not in B neither follow logically from B nor have a probability of 1. K&W's argument against the truth conduciveness of coherence rests on the following premisses:

1. Any non-trivial extension of a belief system is less probable than the original system.
2. There exist non-trivial extensions of belief systems that are more coherent than the original belief system.

Taken together, (1) and (2) entail that more coherence does not imply a higher probability. For by (2), there exist two belief systems, B and B' , where B' is a more coherent non-trivial extension of B . Since B' is more coherent than B and, by (1), less probable than B , it follows that coherence is not truth conducive.

⁸ In a second paper (K&W 1996) they defend their argument against the criticism of Merricks (1995).

How strongly supported are (1) and (2)? Claim (1) is said to require no defence: if the extension is non-trivial, then clearly 'the set of beliefs containing the belief that *p* and the belief that *q* is more likely to contain only true beliefs than the set of beliefs containing both of those beliefs and, additionally, the belief that *r*' (K&W 1994: 130). Claim (2) is said to be supported by the following example:

A detective has gathered a large body of evidence that provides a good basis for pinning a murder on Mr. Dunit. In particular, the detective believes that Dunit had a motive for the murder and that several credible witnesses claim to have seen Dunit do it. However, because the detective also believes that a credible witness claims that she saw Dunit two hundred miles away from the crime scene at the time the murder was committed, her belief set is incoherent (or at least somewhat incoherent). Upon further checking, the detective discovers some good evidence that Dunit has an identical twin whom the witness providing the alibi mistook for Dunit. (K&W 1994: 130-1)

K&W provide the following analysis of the story. Let the original belief system contain the beliefs that (A_1) Dunit had a motive, (A_2) several credible witnesses report that they saw Dunit commit the murder, (A_3) a single credible witness reports that she saw Dunit far away from the crime scene at the time of the murder. Let the extended belief system contain the same beliefs plus the additional beliefs that (A_4) Dunit has an identical twin and (A_5) Dunit did it. The latter system is a non-trivial extension of the former and is more coherent than the former, which is taken to establish (2). K&W draw the moral that 'coherence, per se, is not truth conducive' (1994: 132) and hence that the coherence theory of epistemic justification is untenable.

Clearly, the extended set does hang together better than the original and is therefore more coherent in a pre-systematic sense. This is also what Bonjour's fifth coherence criterion recommends, according to which '[t]he coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system' (1985: 99) The detective's original belief that there is a credible witness claiming that she saw Dunit two hundred miles away from the crime scene can plausibly be seen as confronting the detective with an unexplained anomaly. When the new evidence about the twin arrives, however, the existing anomaly is dissolved (and no new anomalies are thereby introduced).

But do we have to conclude that the extended set in the Dunitz example is less probable than the original set? Let us take a closer look at the example. Superficially, it might seem that the entities that are to be compared as regards their relative probability are the two sets $\langle A_1, A_2, A_3 \rangle$ and $\langle A_1, A_2, A_3, A_4, A_5 \rangle$, and clearly $P(A_1, A_2, A_3, A_4, A_5) < P(A_1, A_2, A_3)$. But this is not to the point. First of all, it is part of the example that the detective *believes these propositions to be true*. What we are to compare are, in fact, not two sets of bare propositions, but two doxastic systems: $\mathbf{S} = \{ \langle BelA_1, A_1 \rangle, \langle BelA_2, A_2 \rangle, \langle BelA_3, A_3 \rangle \}$ and $\mathbf{S}' = \{ \langle BelA_1, A_1 \rangle, \langle BelA_2, A_2 \rangle, \langle BelA_3, A_3 \rangle, \langle BelA_4, A_4 \rangle, \langle BelA_5, A_5 \rangle \}$, where $BelA_i$ says that the subject believes A_i . As we have seen, the principle of total evidence dictates that our probabilities should be based on all available evidence. In particular, this evidence must include facts of belief. Hence what K&W need to establish is that $P(\mathbf{S}') < P(\mathbf{S})$; that is to say, they need to show that $P(A_1, A_2, A_3, A_4, A_5 / BelA_1, BelA_2, BelA_3, BelA_4, BelA_5) < P(A_1, A_2, A_3 / BelA_1, BelA_2, BelA_3)$. But there is no general valid principle saying that extended conjunctions of propositions are less probable *in this conditional sense*. It is true that the larger conjunction will be less probable in an unconditional sense but this loss may conceivably be counterbalanced by the greater strength of the evidence in its support.

Against this it may be objected that, if the detective's beliefs are in fact irrelevant to the truth of the propositions she believes in, we can discard the 'doxastic evidence' when computing the probabilities of the sets in question.⁹ This would allow us to concentrate on the probabilities of the bare propositional sets and conclude that the bigger system is less likely to be true. Yet, it was not part of the example that the detective is completely unreliable in her beliefs, and it would in fact be quite unrealistic to make such a stipulation. One would hope that a detective who is being trusted with a murder case has an impressive track record of solved crimes to support her reliability.

Under the relevant reading, Principle (I) states that, if doxastic system \mathbf{S}' is a non-trivial extension of \mathbf{S} , then $P(\mathbf{S}') < P(\mathbf{S})$. Let us refer to this as the *Doxastic Extension Principle*. An explicit

⁹ This holds under the condition that the beliefs are testimonially independent; see section 7.1.

counter-example to this principle can be found in Appendix A. The example illustrates not only that a non-trivially extended doxastic system need not be less probable than the original system but that it may even be more probable. Thus, by adding new information, we may reach a state of belief that is more likely to be true as a whole.

6.5 Testimonial vs. Explanatory Coherence

It is instructive to compare the present account of testimonial coherence with so-called explanatory coherence. The hallmark of testimonial coherence is that it applies only to ‘sets of supposed facts asserted’, that is to say, propositions that have been reported to be true. This account of coherence includes one form of ‘explanatory coherence’ as a special case, if the explanations in question are such that there is testimonial evidence in their favour. By testimony I mean, as always, any kind of report, including reports from memory, the senses, and other people. Suppose, for instance, that H_1 and H_2 are two propositions of a general hypothetical explanatory character such that H_1 is supported by testimony E_1 and H_2 by testimony E_2 . Then $\{\langle E_1, H_1 \rangle, \langle E_2, H_2 \rangle\}$ qualifies as a testimonial system. If, however, there is no testimonial evidence for a given hypothesis, it cannot be part of a testimonial system. From the testimonial perspective, this means that the hypothesis is useless for the purposes of coherence reasoning.

A variation on the Dunitz example may be used to illustrate these points. Suppose there is testimonial evidence for each of the following propositions: (A_1) Dunitz had a motive, (A_2) several credible witnesses report that they saw Dunitz commit the murder, and (A_3) a single credible witness reports that she saw Dunitz far away from the crime scene at the time of the murder. Consider now the additional propositions (A_4) Dunitz has an identical twin and (A_5) Dunitz did it. If there is testimonial evidence for these propositions as well, the whole system qualifies as a testimonial system. (As the Dunitz example was originally described, there was indeed testimonial evidence for A_4 and A_5 available, since they were supposed to be believed by the detective, and facts of beliefs can be taken as testifying to what is believed.)

Suppose, by means of contrast, that there is no testimonial evidence for either A_4 or A_5 but that they were devised for the sole purpose of jointly explaining A_1 , A_2 , and A_3 . Then, obviously, there is no testimonial system of which A_4 and A_5 can be parts. On the view advanced here, this means that A_4 and A_5 cannot play any role in coherence reasoning. The explanatory coherence theorist disagrees, insisting that the concept of coherence is applicable also to a system containing A_1 – A_5 ; indeed he or she holds that the extended set is more coherent, in the explanatory sense, than the original, smaller set. To be sure, A_4 and A_5 may, when taken together, explain the evidence in the sense of raising its probability. And we may express this fact by saying that there is a positive degree of ‘explanatory coherence’ between the evidence and these two propositions. But what is the great advantage in saying this if it just means that the hypotheses together raise the probability of the evidence? If this is all the explanatory coherentist is saying, then I agree. The problem, of course, is that this is a rather trivial contention.

The explanatory coherence theory becomes dubious if it is combined with a truth conduciveness claim to the effect that explanatory coherence, in the sense that does not require the existence of testimonial evidence for the explanations, is somehow correlated with truth. Consider again the second Dunitz scenario where the only merit of A_4 and A_5 is that they jointly explain the facts of the case. In this case, the probability of A_1 – A_5 will indeed be lower than the probability of A_1 – A_3 , that is to say, $P(A_1, A_2, A_3, A_4, A_5 / E_1, E_2, E_3) < P(A_1, A_2, A_3 / E_1, E_2, E_3)$. A (non-trivially) extended set of propositions is surely less probable than the original set *on the same evidence*. If this is what K&W meant to say, they were right. The point, moreover, has considerable force against strong versions of the explanatory coherence theory. But it does not at all affect the testimonial coherence theory. This shows, once again, that we were charitable to the coherentist in construing her as applying her central concept to testimonial systems.

It should be added, in this connection, that it is often possible to devise several conflicting explanations of a given phenomenon, explanations that cohere equally well with the given evidence. As an alternative to the twin hypothesis we could imagine Dunitz’s ordering one of his partners in crime to dress up like Dunitz so as

to provide him with an alibi. Our unwillingness to lend much weight to 'explanatory coherence' in itself is reflected in the scientific and common sense practice of not accepting an explanation, however well it 'coheres' with the evidence, until it has been confirmed on grounds that are independent. Forensic enquiry is no exception to this general rule. We would not accept the hypothesis that Dunitz committed the crime while his twin was far away, unless we had ascertained by independent means that Dunitz really does have an identical twin, and so on.

Lewis himself sometimes expresses himself in a manner that could lead one to believe that he is advocating a strong theory of explanatory coherence that does not require there to be testimonial evidence for the explanations involved. Thus, after having described his famous example with the 'relatively unreliable witnesses who independently tell the same circumstantial story', he goes on to remark that, of the different hypotheses that may account for the agreement—mainly truth-telling and randomization—'the one hypothesis which itself is congruent with this agreement becomes thereby commensurately well established' (1946: 346). Since, for Lewis, *A* is congruent with *B* just in case *A* and *B* support each other in the probabilistic sense, the statement just quoted is actually a mere tautology. I do not think that Lewis ever took seriously the idea that congruence should be applied to sets involving hypotheses for which there is no testimonial support. Rather, as we have seen, he reserved the application of congruence to sets of supposed facts asserted. Based on such congruence, we may be led to accept this or that hypothesis (e.g. truth-telling) as an explanation of the agreement. But these hypotheses themselves are not part of the system whose coherence is being assessed.

BonJour insists that we 'consider the major role which the idea of explanation plays in the overall concept of coherence' (1985: 98) and that 'the coherence of a system of beliefs is enhanced by the presence of explanatory relations among its members' (ibid.). Conversely, we are told that '[t]he coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system' (ibid.: 99). Nevertheless, none of this is any evidence that BonJour would subscribe to the strong theory of explanatory coherence. He is, to the contrary, quite explicit about

coherence being defined only for belief systems. Since facts of belief count as testimonial evidence, Bonjour's theory is one of testimonial and not explanatory coherence. Paul Thagard has defended a sophisticated version of the strong theory of explanatory coherence. It will be considered in some detail in Chapter 9.

A Negative Answer

7.1 Independence and Individual Credibility

In Chapter 6, I tried to be maximally charitable in my understanding of the coherence theorist's minimal (comparative) truth conduciveness claim. As part of that understanding, I took coherence to be applicable only to testimonial systems, which led me to focus on testimonial truth conduciveness. The resulting account is immune to several objections that have been raised in the literature, above all by Peter Klein and Ted Warfield. What I will try to show in this chapter is that, even under this generous interpretation, the truth conduciveness of coherence would require the satisfaction of certain other conditions as well. Moreover, coherence can be truth conducive at best in a *ceteris paribus* sense. Finally, I will provide reasons that, in my view, strongly indicate that no coherence measure can be truth conducive even in this very weak sense.

My starting point will be the observation that it is implausible to think that coherence is truth conducive in the absence of further conditions: a well-composed novel is usually not true, and yet it may still be highly coherent—perhaps far more so than reality itself or any adequate representation thereof. This raises the question of what the additional prerequisites might be. Before we go on to consider some more substantial conditions, we note that the joint probability of the contents of the reports must not be 0, since nothing could affect the joint probability in that case, with coherence being no exception to this general rule. For the same reason, the joint probability of the contents must not be 1. The satisfaction of these two conditions will be presupposed in the following.

Much of our discussion of independence and credibility in Chapter 2 carries over to the present case. Thus, we may safely

conclude that coherence is not truth conducive if the reports are entirely dependent on each other. In that case, adding additional reports, however coherent with the first, has no effect on the likelihood of what is being said. This means that more coherence among completely dependent reports does not imply a higher posterior probability. As we also noticed, it is implausible to require full independence for coherence to have the desirable effect; intuitively, a tiny influence of the one report on the other does not cancel out the effect of coherence entirely, although it does make that effect less pronounced. Thus, some degree of dependence is compatible with coherence raising the joint probability, but the increase will be less significant than it would have been had the reports been less dependent.

To return to another trivial condition: if, for each proposition A_i , the probability of A_i given report E_i is 1, then the joint probability of the contents of the reports will also be 1, however coherent or incoherent those contents are. If fully reliable reports are admitted, coherence is not truth conducive, but those who have access to such reporters have no use for coherence anyway.

The more important point is that coherence will not be truth conducive unless independent reports have some degree of individual credibility. An individually credible report, as we use that term, lies strictly between full reliability and irrelevance: $1 > P(A/E) > P(A)$.¹ Our discussion in Part I was contingent on the assumption that the reporters say the same thing. In the present context it is important to allow for the possibility that they utter different statements that may be more or less coherent with one another. To establish that coherence, whatever its more precise nature, cannot be truth conducive if the reports are independent but lack individual credibility, we will have to generalize our account of independence. In the interest of focusing on the philosophical points and not on technicalities, I will confine myself to the two-report case.

Suppose a robbery has taken place with Robert being one of the main suspects. There are two witnesses, Helen and Peter, available for questioning. Consider the following propositions:

A_1 = Robert was at the crime scene,

A_2 = Robert had a million in cash the next day,

¹ Note that $1 > P(A/E) > P(A)$ is equivalent to $P(E/A) > P(E/\neg A) > 0$, provided that $1 > P(A) > 0$.

E_1 = Helen says that Robert was at the crime scene,

E_2 = Peter says that Robert had a million in cash the next day.

The underlying intuition is the same as in the simpler case of two reports with the same content. Hence E_1 and E_2 are independent reports on A_1 and A_2 , respectively, just in case there is no *direct* influence between E_1 and E_2 . We check for this by screening off all indirect influences, i.e. influences that pass, as it were, via the supposed facts of A_1 and A_2 . This is accomplished, as before, by conditioning on the supposed facts and their negations. Before, when there was only one supposed fact in play, it sufficed to condition on that single supposed fact and its negation. Where there are two or more supposed facts to consider, there will be more possibilities to take into account.

First, we will make sure that the reports are not directly influenced by other facts than their contents. This involves certifying that Helen's saying that Robert was at the crime scene is not directly influenced by the supposed fact that Robert had lots of cash the next day, as it would be if Helen had inferred that Robert was at the crime scene from her belief that Robert has the money. There might be influences between Robert's supposedly having lots of cash and Helen's saying that Robert was at the crime scene, but these influences will, if the reports are genuinely independent, be mediated via Robert's supposed presence at the crime scene. There is no direct influence between Helen's testimony and Robert's having lots of cash just in case the following hold:

$$\begin{aligned} P(E_1/A_1, A_2) &= P(E_1/A_1) \\ P(E_1/A_1, \neg A_2) &= P(E_1/A_1) \\ P(E_1/\neg A_1, A_2) &= P(E_1/\neg A_1) \\ P(E_1/\neg A_1, \neg A_2) &= P(E_1/\neg A_1) \end{aligned}$$

Similarly, there is no direct dependence between Peter's saying that Robert has lots of cash and Robert being present at the crime scene if and only if the following hold:

$$\begin{aligned} P(E_2/A_1, A_2) &= P(E_2/A_2) \\ P(E_2/\neg A_1, A_2) &= P(E_2/A_2) \\ P(E_2/A_1, \neg A_2) &= P(E_2/\neg A_2) \\ P(E_2/\neg A_1, \neg A_2) &= P(E_2/\neg A_2) \end{aligned}$$

These eight equations together express a sort of evidence–fact independence. The testimonies are independent of facts other than the facts they explicitly report on. We could express this by saying that the testimonies are *focused* on the supposed facts they explicitly report on. This condition is usually satisfied where testimonies are based on observation. In this study, we will be concerned only with testimonies that are focused in this sense. In the discussion of independence in Chapter 2, which was devoted exclusively to the case of different reports on the same supposed fact, the question obviously did not arise as to whether one witness’s report could be influenced by distinct contents of other reports.

Two reports are *testimonially independent* if and only if they are focused on their respective facts in the sense explained above and if, in addition, the following hold:

$$\begin{aligned} P(E_2/A_1, A_2) &= P(E_2/A_1, A_2, E_1) \\ P(E_2/\neg A_1, A_2) &= P(E_2/\neg A_1, A_2, E_1) \\ P(E_2/A_1, \neg A_2) &= P(E_2/A_1, \neg A_2, E_1) \\ P(E_2/\neg A_1, \neg A_2) &= P(E_2/\neg A_1, \neg A_2, E_1) \end{aligned}$$

This definition and the notion of focusing can be generalized to cover any finite number of reports, but for the purposes of this essay it is sufficient to have it defined for the two-report case. Note that independence as defined here is a four-place relation involving two reports and two supposed facts.

Testimonial independence should be sharply distinguished, not only from evidence independence in the sense of $P(E_1/E_2) = P(E_1)$, but also from content independence in the sense of $P(A_1/A_2) = P(A_1)$. In the Robert case, the contents would be independent, in this sense, if Robert’s being at the crime scene would have no effect on the probability that he would have lots of cash the next day. This is clearly implausible in this case. In general, content dependence is perfectly compatible with testimonial independence. On the other hand, such content independence is obviously not compatible with a high degree of coherence.

We are now in a position to show that coherence cannot be truth conducive under the assumptions of full independence and complete unreliability.

Observation: Suppose that the following hold:

- (i) E_1 and E_2 are independent reports on A_1 and A_2 .
- (ii) $P(A_1/E_1) = P(A_1)$ and $P(A_2/E_2) = P(A_2)$.
- (iii) $A_1 \wedge A_2$, $A_1 \wedge \neg A_2$, $\neg A_1 \wedge A_2$, and $\neg A_1 \wedge \neg A_2$ all have non-zero probability.

Then $P(A_1, A_2/E_1, E_2) = P(A_1, A_2)$.

Proof: See Observation 7.1. in Appendix C.

The third condition serves to rule out certain uninteresting limiting cases. Thus, it follows (essentially) already from conditional independence that reports that are completely unreliable regarding their contents, when taken singly, fail to have any effect on the joint probability of those contents, when combined.

We have agreed that a measure C of coherence is truth conducive, in the relevant testimonial sense, if and only if: if $C(A_1, \dots, A_m) > C(B_1, \dots, B_n)$, then $P(A_1, \dots, A_m/E_1, \dots, E_n) > P(B_1, \dots, B_n/F_1, \dots, F_n)$. What happens if the reports are independent and lack individual credibility is that their collective uselessness is added on top, as it were, of their individual uselessness. This is what the foregoing observation says. Thus, under those conditions, checking for testimonial truth conduciveness reduces to assessing whether $C(A_1, \dots, A_m) > C(B_1, \dots, B_n)$ implies $P(A_1, \dots, A_m) > P(B_1, \dots, B_n)$. We owe Peter Klein and Ted Warfield the observation that coherence, whatever its more precise nature, is not truth conducive in this unconditional sense.

7.2 The Need for a *Ceteris Paribus* Clause

In this section, I will discuss yet another condition that must be satisfied for coherence to stand a reasonable chance of being truth conducive. Suppose that we want to evaluate the truth conduciveness of a given measure of coherence. We now find that the propositions reported by one set of witnesses are more coherent, according to the measure, than the propositions reported by another set of witnesses, and also that the conditions we have discussed so far are satisfied; that is, in each scenario the witnesses are individually credible and collectively independent. However, because the witnesses

delivering the more coherent set are less reliable, or more dependent, than the witnesses delivering the less coherent set, the posterior probability turns out to be higher for the latter. If this is correct, then coherence is not truth conducive in the testimonial sense.

Yet it seems just as plausible to take this sort of case as a counter-example not to the truth conduciveness of coherence but to our analysis of truth conduciveness. The factor missing in that account, one could hold, is a *ceteris paribus* clause. Surely, coherence is truth conducive, in the relevant sense, not if it raises the likelihood of truth in an unqualified sense but if it does so *other things being equal*, and in the example the *ceteris paribus* condition is violated since independence and reliability, while being distinct from coherence, are not equal across scenarios. In evaluating truth conduciveness, such factors should be fixed before the degree of coherence is varied. Compare here BonJour's requirement that the conclusion of an argument connecting coherence with truth should be 'that the likelihood that a system of beliefs corresponds to reality varies in proportion to its degree of coherence (and stability) *other things being equal*' (1985: 170, my italics).

But what qualifies as 'other things' here? Can we really assume that reliability and independence belong to the other things? Are there any more 'other things' that need to be fixed? To answer these questions, it will be useful as a preliminary to consider *ceteris paribus* conditions from an abstract point of view.

As Robert L. Frazier notices, other things are equal when there is only one candidate for being an influence on some outcome (Frazier 1995: 114). In excluding other influences, we are trying to isolate the target property. The truth of a statement with a *ceteris paribus* clause depends on what happens when this isolation is achieved. As Frazier (*ibid.*: 119) also observes, this isolation can only be achieved if the properties that have been singled out are capable of independent variation. Otherwise they are not really separable.

What is involved in this capacity for independent variation? In an attempt to answer this question, Frazier asks what it would mean, more precisely, for two properties P_1 and P_2 to be *incapable* of independent variation. That would entail that whenever an object has P_1 to a particular degree, say n , the object also has P_2 to a particular degree, say m , so that 'when one occurs to a particular degree this

is then a certain sign that the other occurs to a particular degree' (ibid.).

While I agree with Frazier that the requirement of independent variation would be violated in this case, I do not think that this sort of situation represents the only type in which it would be violated. Rather, this would be an extreme case. For suppose that whenever an object has P_1 to a particular degree, say n , it has P_2 either to a degree m or to a degree m' . That would, intuitively, amount to lack of independent variation, provided of course that P_2 could take on more than just the two values m and m' before the value of P_1 was fixed. In general the requirement of independent variation is violated if, I take it, fixing the value of the one property imposes limitations on the extent to which the other property can consistently vary.²

How does this apply to coherence? In the present case, we are trying to isolate one property, namely coherence, from other properties that may have an influence on the joint probability of a set of beliefs or testimonies (our 'outcome', to use Frazier's term). Holding the degree of reliability or independence fixed is in perfect compliance with the requirement of independent variation. Fixing these aspects of scenarios does not mean imposing any constraints on what degree of coherence can be consistently attributed. If we learn, for instance, that each witness tells the truth 60 per cent of the time, this information does not put us in a better position than before to eliminate possible coherence values. It is not surprising that coherence and reliability should admit of independent variation: while the former is a property at the content level, the latter concerns the relation between contents and reports.

If we were to learn that the witnesses *always* tell the truth, that would, to be sure, tell us something conclusive about the degree of coherence, for it would enable us to exclude the case of their delivering incompatible testimonies and it would seem therefore that reliability is not independent of coherence after all. Recall, however, that in requiring the reports to be somewhat but not fully credible, we have already ruled out perfectly reliable reports from

² Independent variation, in this sense, is compatible with the instantiation of one property to a particular degree making the instantiation of another property to a particular degree highly likely, so long as that likelihood is less than 1.

consideration. While reliability in general is not independent of coherence, partial reliability is.

In the next two sections, I will enquire whether there are other factors than independence and individual credibility that should plausibly be held fixed in assessing the *ceteris paribus* truth conduciveness of coherence in favourable circumstances.

7.3 Should Specificity (Strength) be Held Fixed?

Tomoji Shogenji has drawn attention to the following problem of coherence and specificity. He asks us to imagine ‘an epistemically ultraconservative agent who only holds a few extremely unspecific beliefs—say, some rocks are heavier than others; some animals sleep sometimes; and someone is humming some tune somewhere’ (1999: 342). In this case it is, Shogenji writes, ‘very likely that her beliefs are all true even though they do not hang together’ (ibid.). Meanwhile ‘a huge collection of highly specific beliefs—such as the entire body of medical science—almost certainly contains errors even though they tightly hang together’ (ibid.). Hence, more coherence does not imply a higher joint likelihood of truth.

This informal argument can be made formally precise using Shogenji’s C_2 -measure of coherence. It can be shown that more C_2 -coherence is not always associated with a higher likelihood of truth. Let T_1 and T_2 be two tautologies, in which case $C_2(T_1, T_2) = 1$. Compare this with the degree of coherence assigned by C_2 to a pair of equivalent but more specific—i.e. less probable—sentences A_1 and A_2 . For instance, if $P(A_1) = P(A_2) = 1/10$, then $C_2(A_1, A_2) = 10$. Hence, the set of A_1 and A_2 comes out as more coherent than the set of T_1 and T_2 . But since T_1 and T_2 are tautologies, the joint probability of the former set is nonetheless lower than that of the latter.³

Now Shogenji, surprisingly, does not want to infer from his specificity argument that coherence, in the sense of C_2 , is not truth conducive. Following him, let us by the ‘total individual strength’ of a set $\{A_1, \dots, A_n\}$ mean $P(A_1) \times \dots \times P(A_n)$. In his example, the total individual strength is very different for the two sets under comparison.

³ Unlike Shogenji’s original specificity argument this example refers to sets of the same size.

From this observation he draws the following moral: '[t]he impact of the beliefs' total individual strength on their truth indicates that we cannot evaluate truth conduciveness of coherence simply by checking whether more coherent beliefs are more likely to be true together than less coherent beliefs' (1999: 342). Rather, 'we need to check whether more coherent beliefs are more likely to be true together than less coherent but *individually just as strong beliefs*' (ibid.).

It follows immediately from the definition of the C_2 -measure that more coherence implies a higher joint probability among sets having the same total individual strength; that is, if $C_2(A_1, \dots, A_n) > C_2(B_1, \dots, B_m)$ and $P(A_1) \times \dots \times P(A_n) = P(B_1) \times \dots \times P(B_m)$, then $P(A_1 \& \dots \& A_n) > P(B_1 \& \dots \& B_m)$. Shogenji thinks that coherence for this reason *is* truth conducive after all. However, as I will argue next, there are two reasons to be displeased with this attempt to save the truth conduciveness of C_2 -coherence.

First, as I have argued on the basis of the work of Lewis and Bonjour, truth conduciveness must be understood in a conditional sense. The relevant question is whether a more coherent set is more likely to be true *given (at least) that the elements are held as beliefs by a given subject*, not whether a more coherent set is more likely to be true in the absence of a believer. Letting $BelA_i$ mean 'Subject S believes that A_i ', our interest concerns the relation between the two conditional probabilities $P(A_1, \dots, A_n / BelA_1, \dots, BelA_n)$ and $P(B_1, \dots, B_m / BelB_1, \dots, BelB_m)$ and not the relation between the corresponding unconditional probabilities. From Shogenji's example with the 'epistemically ultraconservative agent' we may only conclude that the unconditional probability of the more coherent system is higher than the unconditional probability of the less coherent system, provided that strength is held fixed. Nothing can be concluded from the same assumptions about the relation between the conditional probabilities in that particular example.

Nevertheless, Shogenji's tautology example does go through on the testimonial rendering of truth conduciveness. For consider again the formal counter-example. We noted, on the one hand, that $C_2(A_1, A_2)$ is greater than $C_2(T_1, T_2)$ and, on the other hand, that the set of A_1 and A_2 cannot be more probable than the set of T_1 and T_2 since the latter set has probability 1. This is true whether or not these propositions are assumed to be contents of beliefs actually held.

Second, and more seriously, we may ask why the weaker, filtered concept of truth conduciveness should be regarded as more adequate than the stronger, unfiltered one. Shogenji's answer rests on the contention that '[w]e must filter out the effect of the beliefs' total individual strength in evaluating truth conduciveness of *any* epistemic property' (1999: 343). We are told that, unless we keep the strength fixed, no epistemic property whatsoever can be truth conducive. Unfortunately, no argument is offered for this general claim, but only the following 'illustration'. Suppose that we want to evaluate the truth conduciveness not of coherence but of experiential support, and suppose also that Bill has at first, at time t_1 , no experiential support for his belief that 'someone is humming some tune somewhere' which is believed just by hunch. Later, at time t_2 , he acquires a belief in the proposition 'Arnold is humming "Transfigured Night" in his study' on the basis of his auditory experience. Now compare the probability of the less specific belief at t_1 with that of the more specific one at t_2 . It seems that the former probability might well be higher than the latter. For the less specific belief is very likely to be true even in the absence of evidence, whereas identifying Schoenberg's tune from humming is no easy task. This means that sometimes a belief that has experiential support is less likely to be true than a belief with no experiential support at all. But, Shogenji contends, we would not conclude from such examples that experiential support is not truth conducive. Rather, we tend to think that the comparison was unfair to begin with and that it should be restricted to beliefs that have the same antecedent probability.

Convincing as it may be to hold the strength fixed in this particular case, a single example in which strength should be filtered out does not suffice to support the claim that strength should *always* be filtered out. As a matter of fact, it is not difficult to come up with counterexamples to the general claim. Suppose, for instance, that the property we want to evaluate truth conduciveness of is the total individual strength itself. The strength of a set of beliefs is obviously an epistemic property, and yet it would not make sense to filter out the effect of the beliefs' strength in assessing the truth conduciveness of that very property. Hence it is not true that strength should be kept fixed in evaluating the truth conduciveness of *any* epistemic property.

Purged of this assumption, Shogenji's argument for keeping strength fixed in connection with coherence collapses.

Shogenji might retort by excluding strength itself from the class of properties for which it should be filtered out. Is it true, then, that strength should be kept fixed in evaluating the truth conduciveness of any epistemic property, save strength itself? Based on what we have said about the need for a requirement of independent variation in connection with *ceteris paribus* claims, we should require strength to be kept fixed in evaluating the truth conduciveness of epistemic properties that are strength-*independent*, i.e. properties that have nothing to do with strength and hence do not vary with it. Experiential support is plausibly a case in point. Such a requirement is not justifiable, however, in evaluating properties that are strength-*dependent*. Strength can be held fixed when evaluating the truth conduciveness of coherence only if strength and coherence admit of independent variation.

Unfortunately, Shogenji's own coherence measure makes coherence heavily dependent on total individual strength. Indeed the denominator of the defining expression is the total individual strength. To prove the point, note that if strength (i.e. $P(A) \times P(B)$) is not fixed, then $C_2(A, B)$ has no upper bound; it can take on any real value greater than, or equal to, 0. Yet, once we keep the strength fixed, $C_2(A, B)$ does become bounded. For instance, if the strength is assigned a value of $9/10$, then the upper bound is $10/9$. Thus, keeping the strength fixed imposes limitations on the extent to which the degree of coherence can consistently vary, in violation of the requirement of independent variation.⁴

⁴ In Bovens and Olsson (2000: 691) we argued that the prior joint probability should be held fixed in assessing whether coherence is truth conducive. Our argument centred on the following example. Suppose two independent witnesses testify to the effect that a dog died due to witchcraft, which was antecedently very improbable. Suppose also that two other independent witnesses provide us with reports that are not in complete agreement but still not mutually exclusive either, e.g. that the dog died due to cancer and due to pneumonia, respectively. All reporters are supposed to be individually credible to the same degree. Now clearly the former reports are more coherent than the latter reports for the former are in perfect agreement and the latter not. And yet, the posterior (joint) probability that the dog died due to witchcraft is lower than the probability that it died due to the combined effect of cancer and pneumonia. This is so because of the difference in priors. This example was taken to illustrate the thesis that the prior joint probability may influence the posterior joint probability and that we need to keep it fixed on the grounds that we need to keep fixed all factors, other than coherence, that may affect the posterior joint probability. However, the question is whether the prior joint probability is clearly separable from coherence, so that we may legitimately speak of them as two distinct factors. For reasons similar to those advanced here against fixing the strength, I do not think any more that prior joint probability can be isolated from coherence. Fixing the prior may lead to violations of the requirement of independent variation.

Here is an analogy. Suppose that we wanted to investigate the effect of a certain drug on blood pressure by giving the real drug to some patients and placebos to the others. Then, since blood pressure can be affected by the patient's diet, we should require of an appropriate test that the diet be kept fixed, i.e. that all patients be given more or less the same food. After all, the diet is not part of the medication. However, if we wanted to study instead the effect of a certain lifestyle on blood pressure, where 'lifestyle' is taken in a broad sense to include a person's eating habits, there would be no reason why all subjects should be required to be on the same diet. On the contrary, that would be an entirely unmotivated restriction whose implementation would seriously delimit the scope of the study. It is an equally serious limitation, philosophically, to evaluate the truth conduciveness of C_2 -coherence while keeping the strength fixed. This is not the correct test to apply if what we want to find out is whether C_2 -coherence is truth conducive in a *ceteris paribus* sense.

7.4 Should Size be Held Fixed?

Should we hold the size of a testimonial system fixed when we evaluate the truth conduciveness of coherence? That would be reasonable if size were something that could be isolated from coherence in compliance with the requirement of independent variation. Whether this is so turns out to be a rather vexed issue.

As we saw in section 6.1, adding one more report in perfect agreement does not change the degree of coherence as measured by C_1 . This suggests that it may be plausible not to allow the size to vary when assessing the truth conduciveness of C_1 . This is far less plausible for C_2 which is sensitive to the number of agreeing reports. But even in the case of C_1 fixing the size is more problematic than fixing reliability or independence. To be sure, by focusing on information sets of a certain size one has not restricted the assignment of a degree of coherence. On the other hand, when the size is changed as the result of receiving further testimonies, this typically also affects the degree of coherence. By contrast, changing the reliability or independence of the witnesses never brings with it a change in coherence. Size is, in this sense, not completely separable from coherence.

I suspect that this matter cannot be settled convincingly in the abstract but that we need to keep possible applications of coherence reasoning in mind. Once this is done, it seems clear that it would not be informative enough for potential applications to know merely that coherence is truth conducive over sets of the same size. The reason is that we cannot in general assume that the sets we would like to compare will satisfy that requirement. The point is illustrated already by Klein and Warfield's discussion of their Dunitz example, which involved comparing the degree of coherence of one set with the degree of coherence of an extension of that set.

If, as I am inclined to think, size should not be held fixed, C_1 is not truth conducive *ceteris paribus*. Suppose that we have bet money on the outcome of casting a die and, not being present to confirm the result, decide to consult informants who were actually there. Consider two scenarios: in the first scenario, two persons were present to watch the event, each reporting that 'One' came up. In the second scenario, 100 persons were present, out of which 99 report that 'One' came up, and the one remaining that either 'One' or 'Two' did. In the first scenario, there is perfect agreement among the witness statements (and hence a maximum C_1 -value), in the second not (and hence a less-than-maximum C_1 -value). Yet it is clear that other things being equal we would be more confident that 'One' came up in the second, less coherent scenario. Hence, a lower degree of C_1 -coherence may be counterbalanced by a larger number of testimonies. This casts doubt on the doxastic case as well, as it raises the suspicion that a higher degree of C_1 -coherence could easily be outweighed by a sufficient increase in the sheer number of beliefs.

This negative result depends on the particular coherence measure employed. If instead C_2 is used, we cannot conclude that the degree of coherence is greater in the first scenario (with two reports) than in the second (with 100 reports). From what was said in section 6.1, regarding this measure's size-sensitivity, it should be clear that the second scenario will exhibit a much higher degree of C_2 -coherence than the first.

The measures we started out with seemed to capture important intuitions about coherence reasonably well but they turned out, nonetheless, not to be truth conducive *ceteris paribus*. Neither

C_1 nor C_2 is truth conducive *ceteris paribus* under the conditions in question. C_1 , as we just observed, has a problem of size and C_2 has difficulties with strength (specificity). These negative facts raise the question whether there are any coherence measures at all that are truth conducive even in this very weak sense. In order to answer this question we need to take a closer look at the precise nature of the relation between the prior and posterior probability of what is being agreed upon in cases of full agreement. In section 7.5, I will take a closer look at the precise nature of the relation between prior and posterior probability in such cases, focusing on L. Jonathan Cohen's treatment of that issue. That discussion turns out to have direct and, perhaps, surprising consequences for the issue of truth conduciveness. Those consequences are drawn in section 7.6.

7.5 L. J. Cohen on the Influence of the Prior on the Posterior

L. Jonathan Cohen has stressed the extensive dependence of the posterior on the prior in cases of testimonial agreement. To be specific, he thinks that the relationship between the two quantities is one of inverse dependence. Thus he claims, without proof, that '[w]here agreement is relatively improbable (because so many different things might be said) what is agreed is more probably true' (1977: 98). We are told, moreover, that 'when two pieces of circumstantial evidence converge to incriminate the same man, the mathematical probability of his guilt seems to be increased just because there are so many other men that either piece of evidence might have incriminated' (ibid.). This holds, in particular, when the circumstantial evidence consists of corroborating testimonies.⁵ Compare, for example, a case in which Forbes is one of three equally likely suspects with a case in which he is one of just two. Cohen is maintaining, or conjecturing, that Forbes would, on the evidence provided by agreeing testimonies to his disadvantage, be more likely to be guilty in the first case than in the second.

⁵ For the analogy between testimonial corroboration and the convergence of circumstantial evidence, see Cohen (1977: 93).

The two questions I will be focusing on here are, first, whether Cohen's claim is correct as it stands. It will turn out that it is not, not even under a charitable rendering, which leads me to my second question: is there nonetheless a grain of truth in what Cohen says and, if so, where is it to be found?

Before we attempt to evaluate Cohen's claim, we must make sure that we understand it properly. We will continue our practice of referring to the proposition on which the witnesses agree as the hypothesis. Thus, in the Forbes case the hypothesis is H , that Forbes did it. By the prior probability of the hypothesis we mean its probability before the witnesses made their statements. The posterior probability of the hypothesis is the probability after both witnesses have testified.

Cohen's claim is obviously a comparative one. He is not saying that agreement on an antecedently less probable proposition makes its posterior probability very high, only that it makes it higher than it would have been had its prior probability been lower. Superficially, his thesis would therefore appear to be simply the following:

(C1) The lower the prior, the higher the posterior.

In Lewis-type scenarios—whether they be modelled using my model or using the model of Bovens et al. with genuine independence—the posterior indeed increases as the prior approaches zero, but it does so only given a certain level of improbability. Before that level has been reached, the posterior falls with the prior.

This observation suggests that at least the following claim may be true:

(C2) There is a number t such that for all Lewis scenarios, if the prior is below t , the following holds *ceteris paribus*: the lower the prior, the higher the posterior.

But (C2) also fails, and it does so for the following reason.⁶

By differentiation of

$$P(H/E_1, E_2) = \frac{P(R) + P(H)^2 P(U)}{P(R) + P(H)P(U)},$$

⁶ The following remarks are based on my witness model. The same points can be made using the model of Bovens et al.

the posterior takes on its minimum value, in the relevant interval, for

$$P(H) = \frac{\sqrt{P(R)} - P(R)}{1 - P(R)}$$

This number can be interpreted as the surprise threshold, i.e. as the level of improbability at which agreement starts becoming a surprising fact. If we set $P(U) = 9/10$ and accordingly $P(R) = 1/10$, the minimum is reached at $P(H) \approx .24$, as illustrated in Figure 7.1. The absolutely crucial thing here is that *the point at which the posterior takes on its minimum value depends on the probability that the witnesses are reliable*. As can be seen from the equation above, the smaller we choose the probability of reliability (R), the more the valley of the curve will be shifted to the left (referring to Figure 7.1). It should be clear that given any number t we can, by choosing the probability of R small enough, make sure that the minimum will be located to the left of t , thus falsifying (C2). There is no t such that decreasing the prior below t will increase the posterior regardless of the probability of reliability in the scenario in question. That is to say, there is no single surprise threshold valid for all possible scenarios.

What is true, under our independence assumptions (v)–(viii) of Chapter 3, is only the following weaker claim:

(C3) For every Lewis scenario there is a number t such that if the prior is below t , the following holds *ceteris paribus*: the lower the prior, the higher the posterior.




Figure 7.1: The posterior probability of H as a function of its prior probability.

The difference between (C₂) and (C₃) is a matter of swapping quantifiers. The number t is, of course, the surprise threshold of the scenario in question.

I take (C₃) to be the grain of truth in Cohen's thesis about the relation between the prior and the posterior in cases of witness corroboration. We note that our investigation has taken us rather far away from the unqualified claim (C₁) that was our point of departure. (C₃) is much weaker than Cohen's original claim (C₁), in that the scope is restricted in two main ways: (a) to witness scenarios of the Lewis type in which there is uncertainty regarding the reliability of the witnesses, and (b), given such a scenario, to priors below a certain scenario-specific surprise level.

Cohen also thinks that there is dependence in witness settings between the prior and the strength of corroboration. For instance, he maintains that 'the strength of corroboration might seem all the greater if two such severally implausible witnesses agreed independently about a lot of specific details' (1982: 161–2). Cohen is here implying that the strength of corroboration would be greater if the witnesses agreed on 'A brown-haired, one-legged man with a beard did it' than if they agreed merely on 'A man did it'. In the passage just quoted Cohen makes his point in terms of specificity rather than probability, but the connection between these concepts should be clear: if the agreed proposition is relatively specific, the prior will be relatively low. He writes, in the same spirit, that '[w]here the guilt of any arbitrarily selected man is relatively improbable, a particular man's possession of both motive and opportunity is more probably significant' (1977: 98). One way of measuring the significance of a particular instance of agreement would presumably be measuring the resulting strength of corroboration.

Let us now state Cohen's thesis about the effect of lowering the prior on the strength of corroboration in as simple a form as possible:

(SC₁) The lower the prior is, the stronger is the corroboration.

How are we to understand 'strength of corroboration' more precisely? What we are interested in is the degree to which the testimonial facts are favourably relevant to the desired conclusion. It is natural to measure the strength of corroboration of H given testimonies E_1

and E_2 as the difference between the posterior and the prior, i.e. $\text{Corr}(H/E_1, E_2) = P(H/E_1, E_2) - P(H)$.

Just like his conjecture about the posterior, Cohen's thesis about the strength of corroboration turns out to be false in general. However, it is possible to make it true by qualifying it in two different ways: (1) by restricting it to Lewis scenarios, and (2) by adding a *ceteris paribus* clause. Thus, the following is true:

(SC2) For every Lewis scenario: the lower the prior is, the stronger is the corroboration *ceteris paribus*.

Figure 7.2 illustrates the effect on the strength of corroboration of reducing the prior under the same assumptions as in Figure 6.1 (where $P(U) = 9/10$ and $P(R) = 1/10$). Cohen's claim (SC1) about the strength of corroboration, though also false in general, is in a sense closer to the truth than his claim (C1) about the posterior. In order to make (SC1) true, it suffices to focus on Lewis scenarios. This however does not automatically make (C1) true. In order to turn (C1) into a valid statement we need, in addition, to relativize the claim to a scenario-specific surprise level.

The general conclusion is, on the one hand, that Cohen is right in his general thesis about the need for taking the prior into account in the probabilistic analysis of testimonial corroboration; it is a by-product of our examples that the probability of a particular person's




Figure 7.2: The dependence on the prior of the strength of corroboration occasioned by two agreeing testimonies.

guilt, given that he or she has been incriminated by independent witnesses, can show great variation depending essentially only on the prior probability of guilt. On the other hand, our study also shows that Cohen's reasons for his general thesis are incorrect in that he misconceives the nature of this dependence. His assessment of how the posterior varies with the prior is incorrect in general, and when interpreted as a restricted claim about what happens under certain favourable circumstances it is still far from doing justice to the unexpected complexity of testimonial corroboration.

Most of the foregoing discussion of Cohen was carried out, in greater detail, in Olsson (2002*b*). In an interesting reply to that paper, Bovens and his colleagues challenge my interpretation of Cohen as well as my criticism of the thesis I ascribe to him. For one, they contest that Cohen ever intended to claim anything like (C1) and (SC1) or, using their labels, (C) and (C*). Their only argument concerns (C1); they do not provide any reasons for thinking that Cohen might not be committed to (SC1). Hence, I will focus on (C1) and their actually existing argument.

I have taken the following quotation from Cohen as one reason for ascribing to him acceptance of (C1): 'Where agreement is relatively improbable (because so many different things might be said), what is agreed is more probably true' (1977: 98). On my reading, this is elliptical for

(C1') When agreement is relatively improbable (because so many different things might be said), what is agreed is more probably true *than when agreement is relatively probable (because there are fewer things to say)*.

So, for instance, if the witnesses agree on one of five equally likely suspects, that should make what they say more probable than if they agree on one of merely four equally likely suspects. This seemingly natural interpretation led me to focus on the probability of *H* (the agreed proposition) after both witnesses have delivered their testimonies and how that probability varies with the probability of *H* before any witnesses have come forward.

Bovens et al. propose a different way of filling in the ellipsis: when agreement is relatively improbable (because so many different things might be said), what is agreed is more probably true *than when there*

had been only one witness report. In other words, Cohen is construed as making a claim about the conditions under which corroboration takes place, i.e. the conditions under which the addition of one more testimony makes what is being said more probable. Why should he be so interpreted? Because, they contend, his claim is made in a context in which the focus is precisely on those very conditions. On this 'contextual' interpretation, Cohen is construed as claiming:

(C1'') Corroboration takes place when agreement is relatively improbable.

Now this would certainly be a sensible thing to say for someone who believes that corroboration does *not* take place when agreement is *not* relatively improbable. But the problem is that Cohen does not belong to that category. Rather, he thinks, and indeed proves (in his corroboration theorem, see section 2.6.1), that corroboration can happen even if agreement is relatively probable, e.g. when there are only a few things (but more than one thing) that might be said. Bovens and his colleagues solve this problem for their interpretation by suggesting that what Cohen intends is that corroboration takes place *also* when agreement is relatively improbable. But if this is what Cohen really meant, one might wonder why he did not say so in the first place.⁷

Let me move on to the other issue of interpretation which concerns (SC1) and the relevant measure of strength of corroboration. In my evaluation of (SC1), I relied exclusively on the difference

⁷ It is useful, in this connection, to compare what Cohen says in his book (Cohen 1977) with what he maintains in a later paper (Cohen 1982). In the paper he starts out as well by making some remarks concerning the conditions under which corroboration takes place. But he goes on to add: '[i]ndeed the strength of corroboration might seem all the greater if two such severally implausible witnesses agreed independently about a lot of specific details' (161–2). There can be no doubt that this statement goes beyond the mere issue of corroboration conditions; Cohen is here evidently conjecturing a general relationship between the specificity of what is being said (or, what comes to the same thing, its prior probability) and some other quantity, in this case the strength of corroboration. (In my paper, I took Cohen's statement about specificity to be compelling evidence for his commitment to (SC1), and, again, Bovens et al. do not supply any argument to the contrary. The point I want to make here, however, concerns (C1).) Moreover, this statement is made in the very same paragraph in which he discusses the corroboration conditions. My suggestion is that when Cohen contends, in his book, that '[w]here agreement is relatively improbable . . .' he is doing the very same thing: just like in his paper, he is making an additional remark concerning the relationship between the specificity (prior probability) of what is agreed upon and some other quantity. In the paper, the other quantity is the strength of corroboration; in the book, it is the posterior probability.

measure, i.e. on taking strength of corroboration to be given by the difference between the posterior and the prior. Given this measure, I showed that (SC1) is not true in general but holds only under special circumstances. Again, Bovens et al. think that these results are 'fragile', since they are not always stable when instead other measures are employed. Now my own concern was naturally with Cohen's view on the matter. It is true that I did not provide textual support for my particular choice of measure in my article. Yet such textual underpinning can in fact be found, for Cohen (1977: 107) considers measuring 'the probative force of evidence . . . not by the probability of the conclusion on the evidence but by *the difference between this* [i.e. the probability of the conclusion on the evidence] *and the conclusion's prior probability*' (my italics). Clearly, he is here referring to the difference measure as a measure of the probative force of evidence or (what I take to be the same thing) the strength of corroboration. Cohen, as far as I can discover, does not mention anything else that could plausibly be taken as a measure of strength of confirmation. It is interesting, nonetheless, to see whether (SC1) holds up to scrutiny also when other measures are employed, and the competent investigation carried out by Bovens and his colleagues into this matter leaves little else to be desired.

Leaving issues of interpretation behind, let us focus on the evaluation of (C1) and (SC1). Are they true? Bovens et al. agree that they are not valid in general. There remains the issue of how much truth there is to those claims: a heap or just a grain? Bovens et al. make an ingenious proposal for how to settle the matter concerning (C1), asking us to picture a box with the variable r (the probability that the witnesses are fully reliable) on the x -axis and h (the probability of what is being agreed) on the y -axis. They make the further assumption that r is contained in the interval $]0,1[$ while h is contained in the interval $]0,0.5[$. The latter assumption is justified by the fact that we may here focus on the interesting case of antecedently somewhat unlikely propositions. They then raise the question for which pairs $\langle r, h \rangle$ claim (C1) holds. In other words: for which pairs $\langle r, h \rangle$ is it correct to say that if the prior probability of h had been lower, then the posterior probability h^* would have been higher, given the degree of reliability represented by r ? Using my own model of unknown reliability, without genuine independence, they are able

to show that the part of the box for which (C1) holds is considerably larger than the part of the box for which it does not hold. So it would seem that the amount of truth to (C1) is closer to a heap than to a grain.

There are two reasons why I resist this conclusion. First, let us not forget that there are a number of assumptions in play here. Allow me to recapitulate some features of my model to make this clear. One main assumption is that we are not sure how reliable the witnesses are. Otherwise, so I argued in Olsson (2002*b*) and in Chapter 3 of this book, agreement will not be a surprising fact, and nothing like what Cohen conjectures will come forth. What we do know is that they are either reliable or unreliable. This would be a reasonable assumption if we did not know whether or not the witnesses were really present to see anything at the crime scene. If they were, they would be reliable; otherwise not. I made the further simplifying assumption that reliability means perfect reliability (i.e. truth-telling) and unreliability means complete unreliability (i.e. in effect, randomization). Furthermore, my results are relative to a number of different sorts of independence assumptions which, although they arguably normally obtain, are not guaranteed to hold in general. Another substantial premiss is that if the witnesses are unreliable, the chance that they incriminate a given suspect equals the prior probability that this suspect is the real criminal. It is not difficult to imagine circumstances in which this would not be satisfied.

Second, the results Bovens et al. are referring to in this connection are not exactly to the point. As they make clear in their paper, assuming that the witnesses are independent in the standard sense is more natural and also more in line with Cohen's own approach to the subject. Hence, what we should look at is what happens if such independence is added to my model. Then, as Bovens et al. themselves demonstrate, though without commenting on the relevance of this finding to the heap vs. grain issue, there is no significant difference in size between the area of the box for which (C1) holds and the area for which it does not hold. Clearly, the heap argument fails for this reason alone, although it is true that a smaller heap is not necessarily a grain. I will return to Bovens and his colleagues in Chapter 8 where I critically assess the part of their reply that involves the Principle of Indifference.

7.6 The Impossibility of Coherence

What is the relevance of the Cohen debate to the problem of coherence and truth? In the course of the Cohen discussion, two facts were established that will now prove to be of utmost significance to our investigation into the possibility of defining a truth conducive measure of coherence, one that would allow us to conclude from a higher degree of coherence that the contents of the reports are jointly more likely to be true *ceteris paribus*, given their individual credibility and collective independence.

First, the model that was used in the Cohen discussion and that we have encountered already in Chapter 3 in connection with Lewis shows that Cohen was right in his general claim that the posterior varies tremendously with the prior. By choosing the prior differently one can end up with very different posteriors, some that might be sufficient for 'rational acceptance' and some that definitely would not. This is important to our concerns, because it shows that facts of agreement by themselves have little bearing on what we are really interested in, namely the posterior probability of what is being reported.

The extensive dependence of the posterior on the prior has only been established for cases of full agreement in a Lewis-type scenario. But, following the Lewis–BonJour tradition, such scenarios are paradigm cases of coherence. Hence, what holds for scenarios of this type may be expected to hold for coherence in general. The general lesson, then, is that any measure of coherence that is not sensitive to the prior probability of what 'coheres' will be uninformative and hence often useless in coherence estimations of the posterior probability. It is not sufficient for the purpose of such estimation to look at how well supposed facts asserted overlap. An interesting coherence measure will have to be more like Shogenji's C_2 -measure, which is sensitive to the prior, than like C_1 , which lacks such sensitivity.

Second, the Cohen controversy shows that exactly how the posterior varies with the prior depends on the prior probability of reliability. Whether or not we obtain a high posterior, perhaps one sufficient for acceptance, is contingent not only on the prior probability of what the reports say but also on the prior probability of reliability. We can get widely different posteriors depending on how

we choose the latter and, what proves to be absolutely crucial for our purposes, the very *kind* of dependence of the prior on the posterior, i.e. at what level of improbability agreement starts becoming a significant posterior-raising fact, is contingent on the prior probability of reliability. The importance of this observation to our concerns lies in the fact that not only the absolute but also the *relative* height of the posterior, i.e. what is to count as more or less probable conditional on the evidence, will vary with the initial probability of reliability. Now truth conduciveness, in our sense, involves precisely such comparative assessments of posterior probability. Based on this observation one might be led to conjecture that there cannot be an informative measure of coherence that is truth conducive in a Lewis-type scenario. The conjecture turns out to be true:

Impossibility Theorem: There are no informative coherence measures that are truth conducive *ceteris paribus* in a basic Lewis scenario (given independence and individual credibility).

Proof: See Appendix B for definitions of central concepts and a proof of the theorem.

Although the details of the theorem have been deferred to Appendix B, it should be mentioned that the theorem concerns only the case of full agreement and that the probability distributions that are taken into account are assumed to satisfy various conditions, such as independence and individual credibility. While these assumptions may seem restrictive from a formal perspective, they should, based on our discussion in this chapter, in fact be seen as describing fortunate circumstances. Many real-life witness cases will not conform to this model: the witnesses may be useless individually, they may have fudged their story into agreement or their testimonies may show some but not total overlap. We have argued at length that in these cases it is less plausible or even impossible that there could be an interesting measure of coherence that is truth conducive. What the theorem says is that not even under fortunate circumstances can there be any interesting measure of coherence or agreement that is truth conducive in the comparative sense.

Still, it must be conceded that the impossibility theorem is a result about a specific scenario under particular probabilistic assumptions and that it does not by itself disprove the possibility of there being an

informative measure of coherence that is truth conducive in general. Nonetheless, the theorem does become an epistemologically significant fact, I claim, when it is combined with the Lewis–BonJour theory according to which a Lewis-type scenario of full agreement is a paradigm case of coherence. Its status as a paradigm allows inferences to be made from facts of agreement to facts of coherence in general. The works of Lewis and BonJour contain ample examples of this sort of extrapolation. Lewis argued that just as agreement makes what witnesses say likely to be true, so coherence guarantees that our memories can be relied upon. BonJour, by the same token, tried to make likely that the mechanisms that are salient in witness cases also provide the ultimate rationale for our reliance on our cognitively spontaneous beliefs. What we have established is that there is no way to measure extent of coherence in a basic Lewis scenario in a way that is informative. We may conclude, by the witness-belief and witness-memory analogies, that the same is likely to hold for coherence among beliefs and memories. While we have not strictly speaking proved this to hold, it is strongly suggested by the impossibility theorem in combination with the Lewis–BonJour coherence theory.

From the Lewis–BonJour perspective, then, the impossibility theorem strongly suggests that no coherence measure can be both informative and truth conducive. In order for coherence to be truth conducive, more coherence must imply higher posterior probability. But if the *relative* height of the posterior, as a function of the prior, varies with the initial probability of reliability, there is little hope that coherence can be truth conducive. The reason is that coherence, by its very nature, is a property on the level of report contents and is accordingly not allowed to make reference to reliability, which is a mixed-level property that concerns the relation between contents and reports. In short, the posterior is severely underdetermined by facts of coherence, ‘severely’ because considerations of coherence alone do not even allow us to make comparative assessments of the height of the posterior.

Let us see how the view defended here compares with other views. As we have seen, BonJour’s position is that the problem of defining coherence must await a better understanding of many other issues in philosophy, such as the nature of induction and ‘various issues in logic’. While BonJour takes the task of defining coherence to be a

difficult one, there is no evidence suggesting that he would doubt the proposition that we have found so problematic, that coherence is at all definable.

Long before BonJour, however, Ewing made the following interesting remarks:

I think, however, that it is wrong to tie down the advocates of the coherence theory to a precise definition. What they are doing is to describe an ideal that has never yet been completely clarified but is none the less immanent in all our thinking. It would be altogether unreasonable to demand that the moral ideal should be exhaustively defined in a few words, and the same may be true of the ideal of thought. As with the moral ideal, it may well be here that while formulae are helpful, they can provide no complete stereotyped account, and the only adequate approach is one for which there is no space in this book, namely, a study of what our thought can do at its best by means of numerous examples. (1934: 231)

Ewing is here suggesting that it might be impossible to capture the concept of coherence in a formula. But what sort of possibility are we talking about here? On a 'weak' reading, Ewing is saying merely that it would be *practically* impossible, or at least difficult, to state a definition of coherence, as it would require more than 'a few words'. On this interpretation there is little difference in principle between Ewing's view and the view later advocated by BonJour.

However, Ewing's remarks also admit a stronger reading. Thus rendered, he is claiming, more radically, that defining coherence is *logically* impossible, that there is no formula or statement, however long, which could do the job adequately. Thus understood, he is maintaining that there is no systematic account of coherence, and the best we can do is to restrict ourselves to examples.

What we have accomplished in this chapter can be seen as a vindication of Ewing's thesis *under the strong reading*. For it has been argued that it is not only impractical or difficult to define coherence, due to the length of formulas that would be required or to our lack of understanding of crucial notions which such a definition would have to refer to; rather, the impossibility theorem strongly indicates that it is downright logically impossible to devise such a definition.

There is no mystery about this result. In particular, it does not mean that coherence, while being comprehensible to the human intellect, somehow transcends rational definition. It means simply

that the constraints that have been imposed, explicitly or implicitly, on such a definition are jointly incompatible. These constraints include, notoriously, the requirement that a definition of coherence should make that notion, in favourable circumstances, come out as truth conducive *ceteris paribus*. While having coherence imply truth might be too much to ask for, it should at least fall out of a suitable definition that more coherence implies higher probability in a weak *ceteris paribus* sense under normal circumstances (independence, individual credibility). The constraints also include a condition of informativeness: the degree of coherence should give us some information about how high the posterior is, be it only information about its relative height. The whole point, after all, was to use coherence to assess the likelihood of truth in situations of full or partial ignorance about facts of reliability. I have argued that there can be no measure satisfying these requirements. Just as there are no square circles, there is nothing out there that could play the role coherence is supposed to play. The description of that role is itself incoherent. Small wonder there has been so little progress in defining coherence!

The claim that there is literally no such thing as a degree of coherence seems excessive and implausible. Compare a case in which both witnesses say that Forbes committed the crime with a case in which one witness says Forbes did it and another says Jones did it. Surely there is some sense in which the testimonies in the first case are more coherent than those in the latter. Would it not be more accurate for me to say that coherence theorists are mistaken about some of the properties of coherence, not that there is no such thing as coherence?⁸ I resist this conclusion for the following reasons. As for the example, it is natural to assume that we know that the criminal, whoever he was, acted alone. Against that background, the testimonies of the first pair of witnesses are jointly consistent, whereas the testimonies of the second pair of witnesses are not. I agree that the former are more coherent than the latter. But this is an entirely trivial contention which just amounts to saying that a consistent set of testimonies is more coherent than an inconsistent set. The interesting question is surely whether we can compare *consistent* sets—i.e. sets that for all we know may give a correct description of reality—as to coherence.

⁸ This objection was raised by an anonymous referee who also devised the example.

If we can, then this might help us to decide which of these sets is more likely to be true. I claim to have shown that we cannot make any such interesting coherence comparisons between sets that are consistent with our background knowledge. How we express this conclusion is, I believe, not a substantial but a verbal issue. If we choose to formulate it by saying that there is such a thing as a degree of coherence but the coherence theorists are mistaken about one of its properties, then we should hasten to add that the property about which they are mistaken is the very property which made us interested in coherence in the first place—its supposed truth conduciveness. For this reason, I prefer the simpler formulation: there is no such thing as an interesting degree of coherence.

Although I agree with Ewing as to the logical impossibility of defining an interesting measure of coherence, I see no merit in his further contention that ‘it is wrong to tie down the advocates of the coherence theory to a precise definition’. On the contrary, we should urge the advocates of a coherence theory to make explicit what their central concept involves. What happens when this task is taken seriously is that the coherence theory emerges as resting on probabilistically indefensible assumptions. In allowing the coherence theorist to be vague about what coherence means, Ewing is actually legitimizing the very failure he had alerted us to earlier in his book, namely that of robbing coherence of its significance.

It is an ironic fact that the sort of scenario that turns out to be so problematic is the very scenario taken by Lewis and Bonjour to characterize our epistemic predicament as regards our memories and beliefs, the salient feature of such a scenario being our uncertainty as regards the reliability profiles of the reports. In the very sort of scenario they take to be of such fundamental epistemological importance there cannot be a useful coherence measure that would allow us to infer anything interesting, in absolute or comparative terms, about the probability of the contents of those memories or beliefs.

We saw already in Chapter 2, in connection with Huemer’s witness model, that we can get as high a posterior as we wish by adding more agreeing testimonies, even if those testimonies are rather unreliable in themselves. This would suggest that knowledge, unlike what some so-called reliabilists have suggested, is not reducible to ‘belief that has been acquired through a reliable mechanism’.

We can have knowledge without having individually reliable mechanisms, namely if we have agreement. On the basis of the facts that have been presented in this section, in particular the intimate dependence of the posterior on the probability of reliability in such cases of agreement, one might conclude, conversely, that knowledge is not reducible to mere coherence but that facts of reliability must also be allowed to enter into an adequate epistemology. This might lead one to think that a combination of coherence and reliability is what epistemology needs. Indeed, this was the view that emerged from Bovens and Olsson (2002: 147) where it was concluded that reliabilism and coherentism are not opposites but rather complements since the reliability of the process of belief acquisition and the coherence of our belief systems are both factors that influence the posterior probability of the content of our beliefs systems. Based on this observation it was suggested that a reasonable theory of knowledge and justification must be 'ecumenical', ascribing a role to both factors. This proposal was certainly not unreasonable in the context of that paper. Nonetheless, in the light of the results of this section I think it is mistaken, for it presupposes that there is such a thing as a degree of coherence in the first place, one that can be combined with reliability to yield a posterior high enough for rational acceptance. But that claim is thoroughly undermined by the impossibility theorem. The *prima facie* reasonable proposal that we could arrive at a new ecumenical epistemology by combining considerations of coherence and reliability fails since one of the factors to be combined, namely 'coherence', is on closer inspection seen to be non-existent.

Part III

Other Views

This page intentionally left blank

How not to Regain the Truth Connection: A Reply to Bovens and Hartmann

This book can be read as a sustained critique of the idea that coherence could, in any interesting sense, imply truth. In Part II, I argued that coherence is not a coherent concept, that there is no such thing as an autonomous concept of coherence that could play the role such a concept was minimally supposed to play, i.e. that of being indicative of a higher or lower likelihood of truth other things being equal. There is no substantial concept of coherence that is truth conducive *ceteris paribus*, not even under the favourable conditions of independence and individual credibility. This holds in particular for the kind of scenario, characterized by initial uncertainty as regards the reliability profile of the reporters, which coherence theorists have taken to correspond to our initial predicament regarding memories and beliefs.

It remains to consider an objection to this negative conclusion in the form of a competing way of modelling Lewis-type witness scenarios, the main idea being to invoke the Principle of Indifference as a formal representation of our presumed initial ignorance about the reliability profiles of the reporters. To this objection I now turn.

Lewis's problem of initial credibility and the negative result concerning the possibility of specifying a truth conducive measure of coherence have a common source: the underdetermination of the likelihood of truth by facts on the level of report contents and, more specifically, the dependence of that likelihood on the probability of reliability vs. unreliability. The positive proposal in Bovens, Fitelson, Hartmann, and Snyder (2002) can be seen as a suggestion for how

some of this dependence can be got rid of via an invocation of the Principle of Indifference. They made their proposal in the context of Cohen's thesis about the prior–posterior dependence, and so this is where I will start.

In their reply to Olsson (2002*b*), Bovens et al. begin by suggesting a modification of my model of concurring testimony. In the revised model, the witnesses are genuinely independent in the sense explained in Chapter 2. Letting h be the prior probability of the hypothesis H , they calculate its posterior probability given two independent testimonies as follows:

$$P(H/E_1, E_2) = h^* = \frac{(h + \bar{h}r_1)(h + \bar{h}r_2)}{h + \bar{h}r_1r_2}, \quad (1)$$

where r_1 is the probability that the first witness is reliable and r_2 the probability that the second witness is.

Bovens et al. now proceed to argue that the modified model is still unsatisfactory in one respect for it presupposes that we can assign definite prior probabilities to reliability vs. unreliability. Against such models of what they call 'decision-making under risk', they hold that we should instead consider 'decision-making under uncertainty' and plead ignorant as to the definite probabilities of reliability vs. unreliability. The motivation for this strategy is their belief that it better reflects a typical detective's predicament when confronted with witnesses in a criminal case. As they put it, 'in reality, we are often not even capable of making *this* assignment [i.e. of a definite probability assignment to the reliability hypothesis]' (ibid.: 545). Rather, 'the situation is one of decision-making under uncertainty: we have no clue whatsoever about the chance that the witness is reliable or not' (ibid.).¹ One page later they claim that uncertainty is not merely common but even typical of witness scenarios: 'our judgment about the reliability of a witness in a court is typically a judgment neither under certainty nor under risk, but rather under uncertainty' (ibid.: 546).

In order to model the allegedly typical situation, Bovens et al. propose to invoke, as one possible strategy, the Principle of

¹ The terminology stems from decision theory (see, e.g., Luce and Raiffa 1957). Considering the fact that in the case at hand no visible decision is being made it would be more appropriate to speak of *reasoning* under risk vs. *reasoning* under uncertainty.

Indifference. The way they want to employ it is by treating, in effect, all possible probabilities of reliability as equally likely. Technically, they let r_i be a continuous random variable whose values range from 0 (for certainty that witness i is fully unreliable) to 1 (for certainty that the witness i is fully reliable) and proceed to assume a uniform distribution over r_1 and r_2 , calculating the expected value ('average') of h for each value of h in this manner:

$$\langle h^* \rangle = \int_0^1 \int_0^1 \frac{(h + \bar{h}r_1)(h + \bar{h}r_2)}{h + \bar{h}r_1r_2} dr_1 dr_2 \quad (2)$$

At what point does agreement start becoming a surprising fact on this model? Asking this is tantamount to asking for what value of h the posterior takes on its minimum. This will be the level of improbability beneath which a further decrease in the prior occasions a rise in the posterior. To answer that question, Bovens et al. take the derivative of this function with respect to h , set this derivative equal to 0, and solve for $h \in]0, 1[$. What they get is

$$h_{\min} \approx 0.199 \quad (3)$$

This means, as Bovens et al. observe, that Cohen's claim about how the posterior varies with the prior holds true when there are more than five equiprobable suspects, i.e. when h is less than or equal to $1/6$ and thus less than 0.199. Thus, Cohen's claim is, on this model, almost always true and it would appear that I was wrong in claiming there to be but a grain of truth in what he conjectured. The conclusion that Bovens et al. themselves want to draw—that there is a 'heap' of truth in it—would seem more appropriate.

But is the alternative model correct? The first difficulty with this proposal is that Bovens et al. offer no evidence in favour of their claim that in reality we have no clue whatsoever about the chance that a given witness is reliable or not. Let us see what this claim entails. There are many clues that one would expect to have a bearing on the reliability of a given witness's testimony. One such clue would be the witness's own degree of confidence that she has identified the real culprit. This information is usually available, for if we can ask the witness about the identity of the criminal, little seems to prevent us from asking the further question how confident she is that her

identification is correct. Moreover, we would expect a testimony to be more worthy of our trust if the lighting conditions were favourable or if not too much time has elapsed from observation to identification. We are frequently in a position to assess what the conditions were like at the time of observation, and we may well be able to ascertain how much time has passed from observation to identification. Bovens et al. are saying, or implying, that all these bits and pieces of information, which I take to be usually present, are irrelevant to the reliability assessment. One may wonder what the empirical basis of this rather astonishing claim might be.

Let us focus on one of these factors: a witness's own reported confidence that she has got it right. How good an indicator of her actual reliability is this clue? This issue turns out to be a matter of some controversy in experimental psychology.² In a review of thirty-one studies the correlation was found to be close to zero. A later review found a positive but not very impressive correlation.³ More recent research has focused on studying the correlation while varying psychological and statistical factors. For instance, the correlation turns out to be positive when the information-processing conditions are beneficial at the time of observation and identification. It is also positive when the culprit is distinct in the line-up, and, moreover, post-line-up confidence (assessed after the identification) is a stronger predictor of accuracy than pre-line-up confidence (assessed before the confrontation). The received view, based on this later body of research, appears to be that there is a clear correlation between eyewitnesses' confidence and accuracy in easily ascertainable circumstances but that this correlation is not very strong. If we add to this one single clue other clues that are often available, there is, *pace* Bovens and his colleagues, no reason for general scepticism as to the reliability of witnesses and no reason to think that the predicament facing a detective is, in this respect, normally one of complete ignorance. Moreover, if we really do not have a clue as to the reliability, should we not then also admit the possibility that they might be liars rather than presupposing, as Bovens et al. do in

² I here follow Juslin, Olsson, and Winman (1996), which also contains additional references to experimental work in this area.

³ Chapter 15 in Coady's book contains interesting, though not always charitable, criticisms of experimental work that has been taken to indicate a general unreliability of eyewitnesses.

constructing their model, that they are either reliable or produce their reports at random?

Unlike what Bovens and his colleagues think, it is highly problematic to invoke the Principle of Indifference to model realistic witness cases. As we just observed, such cases normally do not involve complete ignorance as regards the reliability of the witnesses. This however need not mean that there cannot be other applications of this model. Did not Lewis contend precisely that we are completely ignorant as regards the specific level of initial credibility pertaining to our memories? What we are supposed to know is only that there is *some* positive credibility. From our supposed ignorance in this respect he concluded that the specific level of individual credibility ‘cannot be assigned’. It turns out that if we use the model of Bovens et al. to model a Lewis scenario the effect is precisely that the degree of individual credibility becomes fixed to a specific number, falsifying Lewis’s claim that ignorance implies unassignability.⁴ Let us see how this works, assuming all probabilities to be non-extreme.

Our point of departure will be decision-making under risk for which Bovens and associates prove the following general theorem (their Theorem 4):

$$P(H/E_1, \dots, E_n) = \frac{h}{h + \bar{h} \prod_{i=1}^n x_i} \quad (4)$$

with the likelihood ratios

$$x_i = \frac{P(E_i/\neg H)}{P(E_i/H)} = \frac{h\bar{r}_i}{h + \bar{h}r_i}$$

for independent witnesses $i = 1, \dots, n$. For $n = 1$ this yields

$$P(H/E_1) = h^{r_1} = \frac{h}{h + \frac{\bar{h}\bar{r}_1}{h + \bar{h}r_1}} \quad (5)$$

which is greater than $h = P(H)$, if r_1 is positive. Disregarding the details, the important thing here is that the individual credibility

⁴ But see below for doubts as to the compatibility of the Principle of Indifference with Lewis’s general philosophy of probability.

depends in this case not only on the prior probability of the hypothesis but also on the probability of reliability.

Equation (5) can now be used as a basis for deploying the Principle of Indifference so as to cancel out the probability of reliability. Following Bovens et al. in assuming a uniform distribution over r_1 , we can calculate the expected value for each value of h :

$$P_1(H/E_1) = \langle h^r \rangle = \int_0^1 \frac{h}{h + \frac{\bar{h}\bar{r}_1}{h+\bar{h}r_1}} dr_1 = \frac{1+h}{2} \quad (6)$$

Proof: See Appendix C for a proof (Observation 8.1).

The individual credibility is positive, i.e. it exceeds h , if, as we have supposed, $h < 1$. Thus invoking the Principle of Indifference is tantamount to assigning to each report a positive degree of credibility that does not depend on the probability of reliability but only on the prior probability of the hypothesis.

Let us contemplate the consequences of the observation just made for Lewis's theory. It was supposed to follow from our ignorance of the individual credibility (or, what comes to the same, the prior probability of reliability) that the degree of such credibility could not be assigned. But our observation shows that this is wrong if we are allowed to use the Principle of Indifference to represent our ignorance. In the model proposed by Bovens et al. the degree of positive individual credibility is as fixed as one could ever wish. With this obstacle removed, the road would be open for coherence to imply a high probability 'by itself' and Lewis's attempted justification of memory could be saved.

Moreover, it would seem that this invocation of the Principle of Indifference would solve our truth conduciveness problem as well. The important thing, from that perspective, is the fact that once the Principle of Indifference has been invoked the point at which the posterior takes on its minimum does not depend any more on the probability of reliability. (We recall that the cause of the impossibility result of the previous chapter was the fact that, whereas the location of the posterior's minimum depends on the probability of reliability, coherence, whatever its more exact nature, is supposed to be reliability-independent.) Referring to equation (3), the

posterior will now take on its minimum in the neighbourhood of 0.199. When we plot the posterior as a function of the prior, this yields one single curve with one single valley and not many different curves with their valleys at different places (depending on the probability of reliability). Once the dependence on the probability of reliability has been eliminated, there is no reason any more to think that there could not be an informative coherence measure that is truth conducive *ceteris paribus* under conditions of individual credibility and independence. It certainly seems possible that such a measure could exist, though it remains an open problem what it would look like.

But how successful is this model, apart from the fact that it undoubtedly succeeds in fixing the degree of initial credibility? A moment's reflection reveals that while it might well solve Lewis's problem of initial credibility, this is not the only convergence parameter whose value has to be assessed for Lewis's vindication of memory to work. The posterior probability of *H* is still dependent on its prior probability. If we have no idea what the prior is, we will not be in a position to estimate the posterior. This is so even if the initial credibility has been fixed.

Yet, even if the problem with the prior can be satisfactorily solved, there is a serious general problem with this model. The Principle of Indifference dictates that, in the absence of any good reason to think that one alternative is more probable than the other, we should assign equal probabilities to all alternatives. But which are the alternatives?

According to Bovens et al. the alternatives concerning whose probabilities we are ignorant are *different probability assignments to the reliability hypothesis*. We do not know the exact probability of reliability and therefore, they reason, we should consider all possible assignments equally probable. But this is only one among several possible ways of construing the alternatives, and it is not even the first that comes to mind. Why not simply say that the alternatives of whose probabilities we are ignorant are simply the *different reliability profiles*, in our case reliability vs. unreliability? In the absence of any reason to think one to be more likely than the other, we might be inclined to assign them equal probability, i.e. a probability of 0.5 each. Let us see where this alternative approach

would lead us. First of all, plugging this assignment into equation (5) yields:

$$P_2(H/E_1) = \frac{h}{h + \frac{0.5\bar{h}h}{h + 0.5\bar{h}}} \quad (7)$$

It can be shown that $P_2(H/E_1) = P_1(H/E_1)$. Whichever way we use the Principle of Indifference, we get the same degree of individual credibility.

So far so good. But let us continue our alternative employment of the Principle of Indifference. Plugging $r_1 = r_2 = 0.5$ into equation (1) yields:

$$P_2(H/E_1, E_2) = \frac{(h + 0.5\bar{h})^2}{h + 0.25\bar{h}} \quad (8)$$

To find the minimum of this function we calculate its derivative with respect to h , set this derivative equal to 0, and solve for $h \in (0, 1)$:

$$\text{The minimum of } f(h) = P_2(H/E_1, E_2) \text{ equals } 1/3. \quad (9)$$

Proof: See Appendix C for a proof (Observation 8.2).

By contrast, the minimum resulting from using the P_1 -function recommended by Bovens et al. is, as we saw, approximately 0.199. So this time we get different results depending on how the Principle of Indifference is invoked. How the posterior depends on the prior will be contingent on what alternatives are to be assigned equal probability. This gives rise to two different prior–posterior curves. No single informative coherence measure can mirror both curves. We would need several different coherence measures, one for each employment of the Principle of Indifference.

It is noteworthy that the disparity between the two uses of the Principle of Indifference concerns the *level* at which the principle is applied; it does not concern what is taken to be a basic alternative. In both applications it was assumed that the basic alternatives as regards the reliability of the reports are reliability (truth-telling) and unreliability (randomization). Rather, the issue is whether the principle should be applied at the basic level of reliability profiles or, as Bovens et al. urge, at the level of probability assignments to those profiles. What constitutes a basic alternative in these cases is a major

issue in itself. We recall here the dispute between Lewis and Bonjour concerning whether to count lying as a basic alternative on a par with reliability and randomness. But even if this matter has been settled, even if we agree about what constitutes the basic alternatives, there will still be a dispute as to how to apply the Principle of Indifference.

But perhaps it should not come as a big surprise that the two different applications of the Principle of Indifference yield distinct results. One could argue that it is one thing to say that this coin will land heads with a 50 per cent chance, quite another to say that, for any values p and q in the interval $[0,1]$, it is as likely that the coin will land heads with probability p as it is that it will land heads with probability q . These predicaments, so the objection continues, are empirically very different, and it is a virtue and not a vice that this distinction is reflected in various uses of the Principle of Indifference. Nevertheless, I think the claim that what we have here are two different states of uncertainty needs to be argued for, not merely stated. To my knowledge, no such argument has been presented by Bovens and his colleagues, or by anyone else for that matter.

It should be mentioned that invoking the Principle of Indifference would, in Lewis's case, lead to a severe clash between his attempted validation of memory and his general philosophy of probability. In the context of the latter, Lewis declared that we always have some evidence to appeal to in assessing probabilities, so that '[c]omplete ignorance of all relevant empirical facts is completely fictitious in the case of any meaningful empirical question' (1946: 309). He added, in the same vein, that '[t]here are no probability problems in the complete absence of empirical data which, directly or indirectly, are indicative of a frequency of past experience' (ibid.). Lewis was decidedly against invoking the Principle of Indifference if, counter to expectation, a case of complete ignorance should be found, saying that 'no reasonable person would apply the Principle of Indifference in such a case' (ibid.). Lewis was here reiterating the view of C. S. Peirce, who insisted that 'when we have no knowledge at all . . . there is no sense in saying that the chance of the totally unknown event is even (for what expresses absolutely no fact has absolutely no meaning), and what ought to be said is that the chance is entirely indefinite' (1878: 179).

Against the backdrop of his separate rejection of the Principle of Indifference, Lewis was right after all in contending that our initial ignorance regarding the degree of individual credibility makes that degree unassignable. Ironically, the only strategy for fixing the individual credibility of memory that would have some chance, however small, of success was one that he had already rejected on other grounds earlier in his book in a general probability context. This points to a severe tension between Lewis's philosophy of probability and his discussion of memory.

Brandt (1954: 89) observes, on distinct but related grounds, that Lewis's general account of probability is irreconcilable with his attempted justification of memory. Lewis's contention that '[t]here are no probability problems in the complete absence of empirical data which, directly or indirectly, are indicative of a frequency in past experience' (Lewis 1946: 309) is made in the context of his general philosophy of probability. In his attempted validation of memory, however, Lewis concedes that there is no empirical evidence to appeal to in support of memory's initial credibility, as any potential evidence would have to be based on memory, the credibility of which is in doubt. It would follow that it is not meaningful to ask whether there is a positive initial credibility of memory as such. But this is evidently not the conclusion Lewis himself draws. Not only does he consider that question meaningful; he even thinks that it should be answered in the positive.

I now take the opportunity to add a brief remark on a new impossibility result by Bovens and Hartmann (2003) which is similar in spirit to the result mentioned in Chapter 7 and proved in Appendix B. Their result came to my knowledge as I was finalizing this book, and I regret that I cannot give a detailed account of their very substantial achievements. Although the book by Bovens and Hartmann with which I am here concerned is listed as published in 2003, it did not in fact appear until mid-2004.

The upshot of their reasoning, too, is that it is impossible to define a general truth-conducive measure of coherence, that is to say, they claim to have solved the problem that was described in Olsson (2002a) as the remaining problem of coherence and truth. That they would make such a claim may come as surprise in the light

of the foregoing discussion. However, the underlying model they use in their impossibility theorem does not employ the Principle of Indifference. Another interesting feature of their book is their proposal for how the coherence theory could be saved from their initial dialectical attack. The main idea is that the impossibility theorem poses a threat only if it is agreed that information sets can be ordered according to their relative coherence in a way that makes all sets comparable. Once comparability is given up, the impossibility result ceases to cast a shadow on the coherence theory, or so they claim. How such a quasi-ordering of coherence can be defined is shown, with much ingenuity and formal sophistication, in the second chapter of their book.

One problem I had with the first two chapters of their book concerns the interpretation of C. I. Lewis, whom Bovens and Hartmann rightly describe as a main advocate of the truth conduciveness of coherence. The impossibility result is based on the assumption that the information sources are reliable to a certain, fixed degree which is not subject to change as more information arrives. This sort of reliability is called 'exogenous' in the book. Lewis, on the other hand, was quite clear about the fact that in the kind of scenario he took interest in the reliability is initially uncertain and subject to subsequent revision. In fact, I cannot think of any coherence theorist who has shown interest in exogenous reliability. What then, I asked myself, is the philosophical relevance of the impossibility theorem? The problem turned out to be one of presentation only; my question was answered in their third chapter where Bovens and Hartmann proceed to take the more complex situation with uncertain or 'endogenous' reliability into account, showing that their impossibility result can be generalized to cover that sort of case as well. Unlike my theorem, theirs does not rely on varying the degree of specificity of information items, and it remains an open question if and how the results are related.

Nonetheless, I remain dissatisfied with the tenor of their discussion of Cartesian scepticism which conveys the impression that the weak (comparative) truth conduciveness claim upon which they focus their attention is all that is needed for the purposes of a coherence theory of justification; and accordingly that the coherence theorists' sole mistake has been to focus unduly on measures of coherence that

impose an ordering, as opposed to a quasi-ordering, on information sets (Bovens and Hartmann 2003: 26–7). In reality, however, weak truth conduciveness does not exhaust the coherence theorist’s conception of truth conduciveness. Bovens and Hartmann fail to mention that C. I. Lewis, for one, was very clear about the need for a more substantial connection between coherence and truth. As we saw in Chapter 3, Lewis thought that we cannot, as a matter of principle, know how reliable our memories are. What we can know is only that they are reliable to some positive degree, though without knowing what that degree is. These considerations led him to urge that, for the purposes of a coherence theory, a high degree of coherence must be taken to imply a high likelihood of truth, regardless of the actual positive degree of reliability of the sources; it is thus insufficient to establish the comparative claim that a higher degree of coherence implies a higher likelihood of truth. Clearly, Bovens and Hartmann’s introduction of quasi-orderings does little in the direction of establishing the more ambitious contention.

Here is a final remark on Bovens and Hartmann’s application of their probabilistic models to Tversky and Kahnemann’s Linda problem as it is described in their third chapter. In the Linda experiment in cognitive psychology, subjects are told that Linda is 31 years old, single, outspoken, and very bright. She studied philosophy as a student and was deeply concerned with issues of discrimination and social justice. Subsequently, the subjects are asked to rank a set of claims according to what they take to be most likely. Two of the claims are (i) that Linda is a bank teller and (ii) that Linda is a bank teller *and* active in the feminist movement. It turns out that a large proportion of subjects will consider the latter claim to be more probable than the former one, in blatant violation of the laws of probability. Bovens and Hartmann’s proposal is that we should assume that claims (i) and (ii) have been reported by some possibly unreliable information sources. This, if reasonable, would open up for the possibility that (ii) can in fact be more probable than (i), more probable, that is, *conditional on the evidence*. If so, the subjects may well have been right after all in their probability assessments. Bovens and Hartmann show that their model can accommodate this possibility under some reasonable reliability assumptions. The strategy, then, is quite similar to our treatment of Klein and Warfield’s detective story where a

similar contrast was made between prior and posterior probabilities (see Chapter 6).

A difficulty facing this proposal is that it was not part of the original Linda example in the first place that claims (i) and (ii) have been reported by some (possibly unreliable) information source. Bovens and Hartmann claim, in response, that '[i]n everyday life, people are typically asked to judge whether a proposition is more or less likely to be true when they have been informed of this proposition by a source (a newspaper, an acquaintance) that may or may not be fully reliable' (85). Therefore, they continue, subjects confronted with the Linda problem are likely to interpret the problem as conforming to the allegedly typical pattern. Still, there are countless cases where we have to compare probabilities without there being any direct evidence on either side. Bovens and Hartmann fail to give any evidence in support of their contention that such situations should be deemed atypical.

Other Coherence Theories

In this chapter, I will discuss, very briefly, the accounts of Nicholas Rescher, Donald Davidson, Keith Lehrer, and Paul Thagard. I will not be able to do justice here to the sophistication of these admirable researchers' work. One reason for this is simply lack of space. But there is also another more substantial reason for avoiding lengthy expositions. While these theories obviously have many merits, they do not, as it turns out, imply much of substance about our problem of whether, and to what extent, the fact that a set of statements is coherent, in the rough sense of its elements 'hanging together', is positively correlated with its joint likelihood of truth. In fact, the traditional notion of coherence that is central to our study plays, as it will turn out, surprisingly minor roles in the works of the first three authors: Rescher, Davidson, and Lehrer. Rescher's central concept is rather plausibility, whereas Davidson is primarily concerned with the external coherence between interpreter and interpretee. According to Lehrer, coherence should be understood in terms of the capacity to answer objections. It can even be doubted whether these theorists are concerned with coherence in the weak sense of 'property definable at the level of contents of some sort of information system'. Paul Thagard's theory will turn out to be more relevant to our concerns and will accordingly be examined at greater length.

9.1 Nicholas Rescher

Central in Rescher's account is the notion of a truth-candidate. A proposition is a truth-candidate if it is potentially true, so that there is something that speaks in its favour. There is obviously a close similarity between Rescher's truth-candidates and Lewis's supposed

facts asserted, upon which my own theorizing was based. Rescher now devotes special attention to the search for what he calls a truth criterion, that is to say, a systematic procedure for selecting from a set of conflicting and even contradictory truth-candidates those which it is rational to accept as truths. His suggestion amounts to dividing the total set of data into maximal consistent subsets and choosing among these subsets. Whereas the traditional concept of coherence as 'mutual support' seems to play a role in the philosophical underpinning of Rescher's theory, the positive theory he eventually devises does not, as far as I can judge, depend in any essential way on that concept. Rather, his favoured method is to base the choice of subset on considerations of relative 'plausibility', a notion that does not bear any direct similarity to coherence.¹ There is certainly much of value in Rescher's theory but, as others have also observed, it is difficult to see its relation to the coherence theory as traditionally conceived.² For this reason, I will refrain from commenting further on Rescher.

9.2 Donald Davidson

Donald Davidson has proposed what he claims to be a coherence theory of truth and knowledge. Davidson's conclusion is that most of a person's beliefs must be true, so that a person's beliefs cannot be massively false. As he also puts it, there is a presumption in favour of the truth of a belief merely because believed. This is clearly a statement to the effect that beliefs have a degree of positive individual credibility. This makes Davidson's theory seem highly relevant in this context. For individual credibility can be seen as a precondition for coherence to have an effect on the likelihood of truth. Yet, as will soon be clear, the notion of coherence as a property of contents of beliefs does not seem to play any major role in Davidson's argument.

Davidson's argument can be stated simply as follows. Consider a case of 'radical interpretation' where understanding has not been secured beforehand. Davidson argues that in such cases we must

¹ As for the relation between plausibility and probability Rescher notes that, owing to certain formal differences having to do mainly with how conjunctions are treated, '[p]lausibility assignments . . . cannot be construed in directly probabilistic terms' (126).

² For this sort of criticism see BonJour (1985: 222–9).

interpret other people's utterances in such a way as to make them in the main true by our own lights. In other words, the interpreter should favour interpretations that as far as possible result in agreement between himself and the interpretee. In support of this claim, Davidson appeals to the Principle of Charity: 'it makes for mutual understanding, and hence for better interpretation, to interpret what the speaker accepts as true when we can' (Davidson 1986: 316). Hence, there is a presumption that what other people believe is true, where truth, again, is judged by the interpreter's lights.

Now Davidson wants to use the foregoing observation as a starting point in an argument showing that a person can know that most of her own beliefs are true, not only by her own lights (presumably she knew that already) but true, as it were, objectively speaking. That she can know this is supposed to be forthcoming once we contemplate the possibility of an omniscient interpreter. The omniscient interpreter, like any other, would have to interpret my beliefs as mostly true by his lights. But since the interpreter is, by assumption, omniscient, what is true by his lights is really and objectively true. Hence, my beliefs are mostly true, not only from my own perspective but objectively speaking.

This may sound too fantastic to be true. One may wonder, for instance, how a *merely* possible omniscient interpreter can serve to guarantee the *actual* truth of most of my beliefs. Moreover, since the Davidsonian strategy applies only if understanding has not already been secured, a person must not make any initial assumptions as to which beliefs the omniscient interpreter will attribute to her. As has been noted by Edward Craig (1990), the net effect is that the person knows that most of her beliefs are true at the expense of not knowing any more what it is that she believes. Davidson's argument trades scepticism about beliefs for scepticism about meaning.

But we may leave the tenability of Davidson's omniscient interpreter argument aside. What is relevant here is that it seems entirely possible to state Davidson's argument without making a single reference to the concept that is central to our study, namely coherence. This impression is quickly verified in the literature where the argument is indeed usually described without that concept being at any point invoked. To be sure, the idea of agreement does play a major role, but this is the agreement between interpreter and interpretee

which the former is supposed to be maximizing or optimizing in the process of making sense of the latter. Such *external* coherence is clearly something other than the *internal* coherence between a person's own beliefs, which is what coherence theorists have traditionally been concerned with.

But perhaps it is part of the optimizing (the term Davidson uses in his later work as opposed to 'maximizing') of external agreement that the interpreter's beliefs are attributed in such a way that a minimum level of internal coherence is attained. Even so, I fail to see how this should imply any substantial connection between coherence and truth. What Davidson's argument can be used for, at best, is to derive a minimum degree of credibility pertaining to beliefs as such. It cannot be used to show that beliefs are credible enough to be regarded as trustworthy. Indeed, it is entirely silent as to the specific level of initial credibility thus derived, and so utilizing it in the context of a coherence justification of belief is problematic in ways that are by now familiar from our discussions of Lewis and Coady, both of whom also employ transcendental arguments in their efforts to establish an initial credibility of the reports they take special interest in. The problem, of course, is that if we do not have any idea of how high the initial credibility is, we will be unable to assess how high a degree of coherence it would take to reach a given threshold of acceptability (assuming, counterfactually, that we could speak sensibly of 'degree of coherence').

9.3 Keith Lehrer

Keith Lehrer is another well-known proponent of what is generally regarded as a coherence theory. Epistemic justification, he tells us, is coherence with a system. In Lehrer (1990), the relevant system was the acceptance system of a given person.³ The acceptance system consists of reports to the effect that the subject accepts this and that. Thus, 'S accepts that *A*' would be a case in point, but not *A* itself. Hence, Lehrer's acceptance system falls under our concept of

³ In the most recent edition of the same book (Lehrer 2000), the relevant system is taken to be the person's evaluation system, which includes not only acceptances but also preferences and 'reasonings'. This more complicated conception will not play any role here.

a testimonial system. An acceptance system can be represented formally as a structure of the following kind: $\{\langle 'S \text{ accepts that } A_1', A_1 \rangle, \dots, \langle 'S \text{ accepts that } A_n', A_n \rangle, \dots \}$. In this respect, Lehrer's theory fits very well into the testimonial framework adopted here. Nonetheless, as I will argue in the following, the concept of coherence that underlies Lehrer's theory is not one of 'mutual support', and so contrary to initial impression his theory is not directly relevant to our project of determining whether there is a substantial connection between coherence and truth.

What is Lehrer's proposal for how to understand the central concept of coherence? His starting point is the fact that we can think of all sorts of objections an imaginative critic may raise to what a person accepts. These objections might be directly incompatible with what that person accepts or they might, while being compatible with the thing accepted, threaten to undermine her reliability in making assessments of the kind in question. For instance, a critic might object to her claim that she sees a tree by suggesting that she is merely hallucinating. That would be an example of the first sort of objection. As an example of the second sort, we might take a case in which the critic replies that the person cannot tell whether she is hallucinating or not. Coherence, and personal justification, is supposed to result when all objections have been met.

Thus, the process of justifying a claim has the character of a game with the objections and answers being the different moves the players can make. Lehrer, fittingly, calls it the *justification game*. If all the objections raised by the critic can be met, then the claimant wins the game. If she wins the game, her original claim coheres with the acceptance system and she is personally or subjectively justified in accepting her original claim; if not, she is not justified in her acceptance (1990: 119). Lehrer is careful to point out that the justification game is only a 'heuristic device for understanding the considerations that make a person justified in accepting something rather than a psychological model of mental processes' (ibid.).⁴

⁴ Lehrer's semi-formal definition of justification runs as follows: *S* is *personally justified* in accepting *p* at *t* if and only if *p* coheres with *S*'s acceptance system at *t*. Further, *p* *coheres with* *S*'s acceptance system at *t* if and only if all objections to *p* are answered or neutralized relative to *S*'s acceptance system at *t*. This raises the question of how the notion of an objection should be understood, and what it might mean that an objection is answered or neutralized relative to an acceptance system. Lehrer defines the notion of an

For all its obvious intuitive appeal, Lehrer's concept of coherence does not seem to have much in common with the traditional concept of mutual support. If one takes it as essential that such a theory make use of a concept of systematic or global coherence, then Lehrer's theory is clearly not a coherence theory. For in Lehrer's view, '[c]oherence . . . is not a global feature of the system' (1997: 31). Rather, what he calls coherence, as we have seen, is a relation between an evaluation system and a proposition. This relation, moreover, 'does not depend on global features of the system' (ibid.).

What reasons, then, are there for calling the relation of meeting objections to a given claim relative to an evaluation system a relation of coherence? As I understand Lehrer, his answer is that it is a relation of 'fitting together with', rather than, say, a relation of 'being inferable from'. In Lehrer (1990), we read the following: '[i]f it is more reasonable for me to accept one of [several] conflicting claims than the other on the basis of my acceptance system, then that claim fits better or coheres better with my acceptance system' (ibid.: 116). He also contends that '[a] belief may be completely justified for a person because of some relation of the belief to a system to which it belongs, the way it coheres with the system, just as a nose may be beautiful because of some relation of the nose to a face, the way it fits with the face' (ibid.: 88). Lehrer is here claiming that a statement's cohering with a system is analogous with a nose's fitting with a face.

However, as I have argued elsewhere (Olsson 1999), this analogy is incompatible with Lehrer's contention that coherence does not depend on global features of a system. For when we say that a nose fits with a face, we mean that combining the two yields a beautiful overall result, so that the nose fits with the face in virtue of the underlying global property of beauty. If cohering is analogous to fitting, as Lehrer proposes, then a statement coheres with a system if combining the two yields a coherent overall result, so that the statement fits with the system in virtue of the underlying global property of coherence. This, again, clashes with Lehrer's declaration

objection as follows: *o* is an *objection* to *p* at *t* if and only if it is more reasonable to accept that *p* on the assumption that *o* is false than on the assumption that *o* is true (relative to *S*'s acceptance system at *t*). An objection *o* to *p* is *answered*, moreover, if and only if *o* is an objection to *p*, but it is more reasonable for *S* to accept *p* than to accept *o* (with the appropriate relativizations). In the 1990 edition of *Theory of Knowledge*, objections were called 'competitors' and answered objections were said to have been 'beaten'.

that coherence does not depend on global features of the system. So Lehrer's relation of coherence with an acceptance system has little to do with coherence properly so called, being more akin to inference.⁵

9.4 Paul Thagard

Paul Thagard's recent model of explanatory coherence bears some resemblance to Rescher's theory. Thagard, too, takes the fundamental problem to be which elements of a given set of conflicting claims to single out as acceptable. While Rescher wants to base the choice on considerations of 'plausibility', Thagard proposes that we use coherence assessments for that purpose. There are many interesting aspects of Thagard's valuable and original theory that I cannot cover here. Rather, I will have to confine myself to what he says about coherence and truth. The main points I would like to make in the following will be (1) that Thagard tends to exaggerate the differences that exist between his neural network model and probabilistic models, (2) that when this is appreciated, many problems from Lewis and Bonjour carry over to the neural net framework, and (3) that Thagard fails, in the end, to present any valid arguments for his claim that there is a substantial connection between coherence and truth.⁶

Let us see how Thagard's theory is supposed to work. What we begin with is, in the epistemological case, a set of propositions. They can cohere (fit together) or 'incohere' (resist fitting together). Coherence relations include relations of explanation and deduction, whereas incoherence relations include different types of incompatibility, such as logical inconsistency. If two propositions cohere, there is a positive constraint between them. If they incohere, this gives rise to a negative constraint. The propositions are to be divided into ones that are accepted and ones that are rejected. A positive constraint between two propositions can be satisfied either by accepting both or by rejecting both. Satisfying a negative constraint means accepting the one proposition while rejecting the other. A coherence problem,

⁵ For a fuller discussion of Lehrer's theory, see Olsson (2003b).

⁶ These considerations are based on Thagard's *Coherence in Thought and Action* from 2000, which is the most recent and comprehensive statement to date of his view on these general issues. All references in this section are to that book.

according to Thagard, consists in dividing a set of propositions into those that are accepted and those that are rejected in such a way that the most constraints are satisfied. Thagard presents several different computational models for solving coherence problems, including a model that is based on neural networks.

Coherence, as Thagard sees it, applies most fundamentally to pairs of propositions (or, more generally, pairs of 'elements') in a network of propositions. Two propositions cohere if there is a preference for accepting both or rejecting both. Similarly, incoherence is, first and foremost, a relation between pairs of propositions that are felt to be incompatible.⁷ Apart from such binary coherence and incoherence relations, he thinks there are three other notions of coherence of potential value: the degree of coherence of an entire set of elements, the degree of coherence of a subset of the elements, and the degree of coherence of a particular element with the rest. Thagard suggests various ways for how to measure these quantities. I will return to this below.

How acceptability depends on coherence (in the typical case of a network contains evidential data as well as hypotheses) is codified in Thagard's 'principles of explanatory coherence':

Principle E1 (Symmetry): Explanatory coherence is a symmetric relation, unlike, say, conditional probability. That is, two propositions *A* and *B* coherence with each other equally.

Principle E2 (Explanation): (a) A hypothesis coheres with what it explains, which can either be evidence or another hypothesis. (b) Hypotheses that together explain some other proposition cohere with each other. (c) The more hypotheses it takes to explain something, the lower the degree of coherence.

Principle E3 (Analogy): Similar hypotheses that explain similar pieces of evidence cohere.

Principle E4 (Data Priority): Propositions that describe the results of observation have a degree of acceptability on their own.

⁷ Thagard's basic theory does not handle incompatibilities that involve more than two elements in an essential way, e.g. 'Al is taller than Bob', 'Bob is taller than Cary', and 'Cary is taller than Al'. However, he reports the existence of computational methods for converting constraint satisfaction problems whose constraints involve more than two elements into binary problems (19), concluding that his characterization of coherence 'suffices in principle for dealing with more complex coherence problems with nonbinary constraints' (ibid.).

Principle E5 (Contradiction): Contradictory propositions are incoherent with each other.

Principle E6 (Competition): If A and B both explain a proposition, and if A and B are not explanatorily connected, then A and B are incoherent with each other (A and B are explanatorily connected if one explains the other or if together they explain something).

Principle E7 (Acceptance): The acceptability of a proposition in a system of propositions depends on its coherence with them.

How different is Thagard's explanationist framework from the probabilistic setting adopted in this book? In my view, Thagard tends to overemphasize the difference between explanatory and probabilistic coherence and in particular the supposed advantages of the former over the latter. Thus, he claims that 'the explanationist approach sees no reason to use probability theory to model degrees of belief', his reason being that probability theory, while being an 'immensely valuable tool for making statistical inferences about patterns or frequencies in the world', 'is not the appropriate mathematics for understanding human inference in general' (249–50).

The upshot of Thagard's subsequent detailed comparison of the two frameworks is that if computational issues are disregarded we cannot say which framework is better: it is non-trivial but possible, at least in principle, to translate between the frameworks and 'it is an open question whether explanationist or probabilist accounts are superior' (271). Nevertheless, considerations of efficiency speak in favour of the former: 'probabilism might reign supreme in the epistemology of eternal beings... [b]ut explanationism survives in epistemology for the rest of us' (272–3).

How different is Thagard's account of coherence from the conception that I have tried to shed light on? As I understand coherence, it is a property of a testimonial system. A testimonial system, we recall, is a set of pairs $\{\langle A_1, E_1 \rangle, \dots, \langle A_n, E_n \rangle\}$ where E_i constitutes testimonial evidence for A_i . The evidence can, for instance, come in the form of testimony from other people, from memory, or from the senses. In the Lewis–BonJour tradition, as I reconstruct it, coherence is applied only to structures of this general kind. Lewis, for example, tended to focus on coherence among a person's own memories. It is

true that such coherence can raise the probability of other propositions of a purely hypothetical nature, e.g. the hypothesis that the evidence is reliably reported. But this is quite possible without any assessment of the coherence of the hypothesis with the evidence ever taking place. Of course, we could say, in such cases, that the hypothesis is coherent with the data, and Lewis did sometimes adopt this manner of speaking. But I fail to see the point in so doing. Thagard's conception is different from the Lewis–BonJour theory since, in his theory, there are no constraints on what sort of proposition can figure in a coherence problem and hence no restriction on what sets of propositions can 'cohere'. Sets of propositions in a network will not in general be describable as testimonial systems. Typically, some propositions will have the status of evidence and others the status of (mere) hypotheses that were devised only to explain the evidence.

Let us now turn to the problem of coherence and truth. Thagard concedes that 'a major problem arises when we try to justify coherence-based inference with respect to a correspondence theory of truth' (78). The problem is 'most acute' (*ibid.*) for pure coherence theories that do not assign priority to observational elements. The difficulty is less pressing for Thagard's own theory as it 'gives priority (but not guaranteed acceptance) to elements representing the result of observation and experiment' (78). Thagard spends some time elucidating how such priority can be accomplished in various implementations of his theory, the general idea being that of spreading activation first to units representing observational elements and only later to other units whose activation depends heavily on that of the elements in the first class. He proceeds:

Therefore, if we assume with the correspondence theory of truth that observation and experiment involve in part causal interaction with the world, we can have some confidence that the hypotheses adopted on the basis of explanatory coherence also correspond to the world and are not mere mental contrivances that are only internally coherent.

Leaving this unorthodox understanding of the correspondence theory of truth aside, the suggestion is that hypotheses that maximize coherence in Thagard's sense are at least somewhat likely to be true, provided that the observational propositions that enter as input have

some degree of initial reliability, in the sense that they 'are known to make a relatively reliable contribution to solutions of the kind of problems at hand' (73).

Although this is an appealing suggestion, based on what has been argued before in connection with the requirement of individual credibility, there are still two main problems with Thagard's approach.

First, while maximizing coherence may be confidence-raising, provided there are observational propositions and they have an initial degree of credibility, there is no guarantee that there will be any such observational propositions. There is no requirement on a coherence problem, as Thagard defines it, that some of the propositions involved should record observations. The obvious solution to this problem is to add a requirement to the effect that every coherence problem should involve some propositions of an observational nature along the lines of Bonjour's Observation Requirement.

Second, and more seriously, even if some observational propositions are taken into account, there is intuitively no guarantee that coherence will be a significant fact. While coherence among independent pieces of evidence may well be, and often is, confidence-raising, such coherence among dependent items is simply worthless. Since there is no mechanism in Thagard's networks for representing independence, he cannot make this distinction formally. In practice it seems that Thagard tacitly assumes that at least some items of evidence that enter into the networks are independent in the sense of not deriving from one and the same source.⁸

Given these two observations, it is clear that there is no plausibility to the claim that maximization of explanatory coherence should produce true conclusions *in general*. Thagard's own principle of data priority is not sufficient to support any such conclusion. It would at least have to be added that there actually is observational input entering the network and, moreover, that the different pieces of evidence that provide the data are in fact testimonially independent.

Interestingly, the truth conduciveness claim that Thagard finally arrives at is different from his first thesis that 'we can have some confidence that the hypotheses adopted on the basis of explanatory

⁸ See, for example, his representation of the Holmes-Watson example (267).

coherence also correspond to the world and are not mere mental contrivances that are only internally coherent' (79). For he notes that this thesis does not rhyme well with his other claim, made earlier in the same book, that '[i]t may turn out at a particular time that coherence is maximized by accepting a set A that is inconsistent' (75). The latter is indeed a formal consequence of Thagard's theory. Thagard does indicate one way of blocking this consequence through a modification of his coherence conditions (73). At the same time, he makes it clear that he is not in favour of any such change because it would be dissonant with what he takes to be scientific practice: 'Quantum theory and general relativity may be incompatible, but it would be folly given their independent evidential support, to suppose that one must be rejected' (74). So, presumably, the theory that maximizes coherence in a given case can be one that is inconsistent and therefore necessarily false.

In an effort to resolve this apparent internal difficulty, Thagard suggests that we should look at the matter over a longer period of time:

Given a correspondence theory of truth and the consistency of the world, a contradictory set of propositions cannot all be true. But no one ever suggested that coherence methods guarantee the avoidance of falsehood. All that we can expect of epistemic coherence is that it is generally reliable in accepting the true and rejecting the false . . . Temporary tolerance of contradictions may be a useful strategy in accomplishing the long-term aim of accepting many true propositions and few false ones. Hence there is no incompatibility between my account of epistemic coherence and a correspondence theory of truth. (79–80)

Thagard is thus led to reject the simple inference from 'being acceptable on the basis of explanatory coherence' to 'being somewhat likely to be true'. Instead, he now maintains that coherence implies truth in the sense that coherence-based acceptance is conducive to the long-term goal of 'accepting many true propositions and few false ones'. This shift may be sufficient to remove the immediate threat posed by the difficulty involving inconsistency.

More important than the incompatibility problem, however, is the question whether coherence really is conducive to truth in Thagard's new sense. In support of this contention, Thagard appeals to the

perceived pragmatic success of scientific thinking and to his conviction that such thinking is based on the method of explanatory coherence:

Scientific thinking based on explanatory and analogical coherence has produced theories with substantial technological application, intersubjective agreement, and cumulativity. Our visual systems are subject to occasional illusions, but these are rare compared with the great preponderance of visual interpretations that enable us successfully to interact with the world. Not surprisingly, there is no foundational justification of coherentism, only the coherentist justification that coherentist principles fit well with what we believe and what we do. (79–80)

As I understand Thagard, he is reasoning as follows. We all think that scientific thinking is pragmatically highly successful and, moreover, that it is based on the method of explanatory coherence. Now given its pragmatic success, it would be surprising if explanatory coherence were not conducive to truth. Thus, considerations of explanatory coherence dictate the hypothesis that such coherence is conducive to truth (in Thagard's sense) as the only plausible explanation of the granted pragmatic success of science.

But at this point it might be objected that what we may appeal to here is only the *perceived* pragmatic success of science, not its actual pragmatic success. Else we would be begging the question against various deception hypotheses. Systematic deception—whether it is practised by a Cartesian demon or by students of brains in vats—would be just as plausible an explanation of such perceived success as truth would. As far as I can see, there is no indication that Thagard would want to follow C. I. Lewis in ruling out such possibilities beforehand. Nor does he try to show that they are less coherent with our beliefs than the truth hypothesis is.⁹

Thagard does consider 'idealism', which he takes to be the view that there is no external world but that everything is in the mind, as a possible hypothesis of perceived pragmatic success. However, he

⁹ The suggestion is to justify the truth conduciveness of the coherence method with reference to that method itself. This 'inference to the best explanation' would be viciously circular, if the coherence method relied for its justification on its conduciveness to truth. The latter claim is however not something Thagard would subscribe to. Rather, he sees the rationale for the coherence method in the way it systematizes various sorts of reasoning—scientific and commonsensical—which we perceive to be sound on pre-systematic grounds.

believes that some aspects of observation are difficult to explain within a purely coherentist, idealist perspective (88). For instance, there is a sense in which people cannot observe what they want, since most sensory experience is beyond conscious control. Different people in the same situation report very similar experiences. Thagard also mentions as relevant that observations of rocks, fossils, and archaeological sites suggest that the planet Earth has existed for billions of years, but that humans came much later into the picture. Yet, what we can take as data here are only our *perceptions* of the factual status of these claims; we cannot take the claims themselves as data. A sufficiently systematic deception will be phenomenologically indistinguishable from realism, and there is no reason to think that coherence considerations could favour the one hypothesis over the other from a neutral point of view.

Thagard maintains that we have no reason to believe that we are systematically deceived. In support of this view, he maintains that deception hypotheses are lacking in simplicity and hence also in explanatory coherence, as Thagard understands the latter notion.¹⁰ He thinks that ‘materialism’, which he understands as the hypothesis that there is an external world, does not share this deficiency. I agree with Thagard that we have no independent reason to believe that we are deceived. But I do not believe that this result is forthcoming by giving deception a neutral hearing and comparing the different possibilities with respect to their explanatory or other coherence. The main problem is that, if we are really to give deception the benefit of the doubt, then we should consider it from a doxastic position that is in fact not informative enough to allow meaningful comparison of the alternatives. In particular, that position would not be informative enough to allow meaningful comparison of explanatory coherence. In Chapter 10, I will argue that we should not give the sceptic a neutral hearing in the first place.

Although I am sceptical concerning Thagard’s in my view too optimistic assessments of the relation between coherence and truth, much of what he says in other respects is in conformity with the account of coherence offered in the present book. One interesting point of agreement concerns our shared pessimism as to

¹⁰ Cf. Thagard’s discussion of the God hypothesis (89).

the possibility of measuring coherence. After having raised the issue, Thagard makes the following declaration:

It would be desirable to define, within the abstract model of coherence as constraint satisfaction, a measure of the degree of coherence of a particular element [with the rest] or of a subset of elements, but it is not clear how to do so. Such coherence is highly nonlinear, since the coherence of an element depends on the coherence of all the elements that constrain it, including elements with which it competes. The coherence of a set of elements is not simply the sum of the weights of the constraints satisfied by accepting them, but depends also on the comparative degree of constraint satisfaction of other elements that negatively constrain them. (39)

Thagard goes on to say that this observation casts doubts on the very possibility of quantifying statements such as ‘Darwin’s theory of evolution is more coherent than creationism’. While Thagard raises these concerns on grounds that are peculiar to his network conception of coherence, my own reason for doubt was that there is no way to define a useful and informative concept of degree of coherence such that more coherence implies a higher likelihood of truth *ceteris paribus*. Despite these differences, I very much share Thagard’s conclusion.

These remarks are undoubtedly too impressionistic to be conclusive. Still, they give some indication, I hope, that Thagard must agree with our previous conclusion that coherence *by itself* is not correlated with a high likelihood of truth, if only for the simple reason that the role of independence and individual credibility cannot be disregarded. All in all, there is little in Thagard’s argumentation that would suggest that coherence should be truth conducive, be it only in a weak *ceteris paribus* sense.

Part IV

Scepticism and Incoherence

This page intentionally left blank

Pragmatism, Doubt, and the Role of Incoherence

10.1 Cartesian Scepticism

Although a full treatment of scepticism falls outside this essay, which is devoted primarily to the problem of coherence and truth, it is unsatisfactory to dismiss coherence theories without somehow filling the void that is thereby created. As we saw in Chapter 3, C. I. Lewis defends his assumption of an initial credibility of memory on essentially pragmatic grounds. One may wonder whether he should not have invoked pragmatic considerations much earlier in his epistemology. The main issue in this chapter will be whether a more consistent form of pragmatism is defensible. I will argue that it is. My starting point will be Cartesian scepticism and systematic deception. I will try to make clear as we progress how the discussion is relevant to the projects of Lewis, BonJour, and Coady.

The radical sceptic argues that we do not have any knowledge of ordinary things since we do not know that we are not systematically deceived. For instance, I do not know that I am a normal person sitting in front of my computer since, she claims, I do not know that I am not a brain in a vat artificially stimulated to have just those impressions. Recent years have seen the development of a number of sophisticated and precise responses to this sort of scepticism. Contextualism, as advocated by David Lewis and others, belongs to this category as does the relevant-alternatives theory of Fred Dretske and his followers. By contrast, so-called ‘pragmatist’ responses to scepticism seem to have left no visible trace in the

contemporary analytical debate.¹ Anyone familiar with American pragmatism will know that this movement was driven at least partly by anti-sceptical sentiments, and that much effort was spent on explaining what is wrong with radical doubt. One can speculate about the causes of its present lack of influence. The pragmatists, unlike most contemporary researchers, did not focus directly on the sort of scepticism that worries about sceptical alternatives, like brains in vats. And it is not evident how their criticisms of other sorts of scepticism translate into an objection against scepticism of this nature. Peirce's rejection of the method of Cartesian doubt is a case in point. Another possible cause may be that the new approaches are often clearly stated in modern analytical style, whereas the understanding of the pragmatists usually requires a greater creative effort on the part of the reader.

In this chapter, I will be focusing on two main candidates for a pragmatic approach to radical doubt. One is often ascribed to William James and amounts to applying his 'wager' argument against religious agnosticism, as it is stated in his article 'The Will to Believe', in this new context. The second proposal, which can be attributed to Peirce, consists in a direct attack on one of the sceptic's main premisses: that we do not know that we are not systematically deceived.

My main theses will be, first, that the Jamesian proposal is incoherent: the argument against religious agnosticism is not applicable to the issue of radical scepticism. Second, I will attempt to show that the Peircean response contains the core of a plausible pragmatist reply to scepticism, although his account is in one important respect incomplete. Finally, I will argue that the Peircean reply has some advantages over the other main contemporary responses, as they have been articulated in the literature. The disposition is dialectical rather than chronological. After having presented the sceptical argument, I turn to James. The rest of the chapter is devoted to the Peircean type of response.

I should point out that the Peirce discussion has interesting implications for the problem of assigning an epistemic role to the concept of

¹ In a meticulous survey of recent work on radical scepticism, with well over a hundred references, pragmatism is not mentioned as an alternative position (Pritchard 2002). Vaguely related ideas are discussed under the heading of neo-Moorean responses. For more on the relation between those responses and pragmatism, see section 10.6 below.

coherence. A central part will be devoted to the legitimate grounds for doubt. Peirce's central idea in that regard is that doubt is always induced by some sort of incoherence. Taking his rather restrictive view on what counts as incoherence as my point of departure, I proceed to isolate a more extensive class of types of cognitive dissonance. The proposal is that while coherence may lack the positive role many have assigned to it, mainly due to the lack of a correlation with likelihood of truth, incoherence plays an important negative role in our enquiries.

The problem of scepticism can be posed in the form of a paradox, consisting in the joint incompatibility of three claims, each of which appears, taken individually, to be acceptable. We may let *SH* refer to any sceptical hypothesis, such as the hypothesis that I am a brain in a vat the experiences of which are the effect of artificial stimulation of nerve cells. Let *O* be some common sense proposition which I claim to know. *O* should be chosen so as to entail the falsity of the sceptical hypothesis under consideration. For instance, we could let *O* be the proposition that I have two hands. The three incompatible claims are then the following:

- (1) I know *O*.
- (2) I do not know not-*SH*.
- (3) If I do not know not-*SH*, then I do not know *O*.

These three propositions are jointly incompatible and yet each seems at least *prima facie* acceptable. The first premiss is acceptable since *O* is supposed to be something we think we know. The second claim is plausible because *SH* is the very sort of proposition which one seems unable, in principle, to know: it concerns a scenario, such as a brain-in-a-vat scenario, which is phenomenologically indistinguishable from everyday life. It is logically possible that what we perceive is entirely an illusion, and there does not seem to be any reason to prefer the normal view to the illusion hypothesis; they are entirely symmetrical.² Finally, the third premiss also seems true. If I know I have two hands, then I know I am not a brain in a vat.

² Franklin (1991: 309) calls the symmetry argument 'the best argument for scepticism'.

The sceptic takes the plausibility of (2) and (3) as a reason for rejecting (1), that is to say, she is arguing as follows:

(S1) I do not know not-*SH*.

(S2) If I do not know not-*SH*, then I do not know *O*.

Hence:

(S3) I do not know *O*.

For example,

(S1 \star) I do not know that I am not a brain in a vat.

(S2 \star) If I do not know that I am not a brain in a vat, then I do not know that I have two hands.

Hence:

(S3 \star) I do not know that I have two hands.

The same argument is equally effective against any common sense knowledge claim, and so, the sceptic concludes, we hardly know anything at all.

10.2 Jamesian Wagering

Let us add the following premiss to the sceptical argument:

(S4) If I do not know *O*, then I should stop believing *O*.

From (S3) and (S4) we may derive:

(S5) I should stop believing *O*.

From (S1), (S2), and (S4) we may conclude, for instance, that I should stop believing that I have two hands. The same reasoning applies to any one of my common sense beliefs. Hence,

(*UC*) I should give up all my present beliefs.

I will refer to (*UC*) as the ‘uncertainty consequence’ of the sceptical argument (supplemented with premiss (S4)).

But is the sceptic really committed to (S4)? Why should we attribute to her the view that we are not entitled to believe what we agree we do not know? Well, if you agree that you do not know where you parked

your car, it would surely be strange for you to hold on to your belief ('belief' here, as always, in the sense of 'full conviction') that you parked it behind the supermarket. Moreover, it is difficult to see the relevance of the sceptic's argument, were it not supposed to have (*UC*) as a consequence. Surely, she does not only want to suggest that we do not know many of the things we think we know, but also that we must give up all those beliefs which, in her view, have turned out not to be known, so as to treat everything as uncertain. Thus, the potential importance of the sceptical argument lies in (*UC*) which is not forthcoming unless (*S4*) is added to the sceptic's premisses. Without (*S4*) and (*UC*), the 'problem of scepticism' becomes an empty exercise without any clear bearing on human enquiry.

H. O. Mounce (1997: 100) notices that the sceptical argument seems to have the uncertainty consequence:

Thus, let us grant the perspective adopted by the sceptic... The point might be expressed by saying that we have no absolute certainty. But what consequences are supposed to follow from that point? For example are we supposed in practice to face the future with no expectations whatever, treating everything as uncertain?

As we will see, Mounce's response to this is that treating everything as uncertain is an irrational strategy from a practical point of view. The demonstration that the uncertainty consequence is practically absurd consists in the appeal to a wager argument, in analogy with James's famous defence of religious belief in 'The Will to Believe'. (James, in turn, was heavily influenced by Pascal's wager argument in his *Pensées*.) It is practically better to stick to one's normal beliefs, as opposed to giving everything up.

Supposing for the sake of the argument that treating everything as uncertain is irrational, what would follow? Only that we have reasons to think that one of the premisses—(*S1*), (*S2*), or (*S4*)—leading up to (*UC*) must be rejected and that this has to be done on practical grounds. The argument can be seen as a practical *reductio ad absurdum* of the sceptic's concept of knowledge. Mounce himself is a case in point. In his view, the sceptical paradox arises because knowledge is incorrectly construed as an absolute notion. This is supposedly a misconception because 'our knowledge is always relative to normal conditions that pass beyond our knowledge' (99). Although Mounce

is not explicit on this point, it seems that he would reject (S₂): we can know normal things without knowing that sceptical hypotheses are false, the reason being that going from one of these knowledge claims to the other means shifting perspective.³ Before turning to the details of Mounce's argument, it will be helpful to recapitulate the central ideas of James's paper.⁴

Pivotal in James's discussion is an illuminating taxonomy of different kinds of options, i.e. of decisions involving two hypotheses. Such an option is *living* if both hypotheses are live ones, i.e. if they are not already excluded but make some appeal, however small, to the person's beliefs. An option is *momentous* if it is important, e.g. because something of great value can be won and the opportunity is unique. Finally, an option is *forced* if it is, in a sense, not avoidable. If the two hypotheses form a 'complete logical disjunction' in the sense of being exhaustive, the option is forced. For instance, 'Either accept *H* or not' is a forced option in this sense: you cannot avoid choosing one of the alternatives; whereas the option 'Either love me or hate me', to take one of James's own examples, can be avoided by doing neither. But even an incomplete logical disjunction can be forced, if avoiding choosing one of the hypotheses is indistinguishable, so far as the consequences are concerned, from choosing one of the hypotheses. An option is said to be *genuine* if it has all these characteristics, i.e. if it is living, momentous, and forced.

James advances two theses about genuine options. His first thesis runs as follows (II):

Our passional nature not only lawfully may, but must, decide an option between propositions, whenever it is a genuine option that cannot by its nature be decided on intellectual grounds; for to say, under such circumstances, 'Do not decide, but leave the question open,' is itself a passionate decision,—just like deciding yes or no,—and is attended with the same risk of losing the truth.

What James is opposing is the position that suspending judgement is the uniquely rational or 'intellectual' strategy in cases of the sort just described.

³ Thus interpreted, pragmatism becomes difficult to separate from contextualism. For more on contextualism, see section 10.6 below.

⁴ Unless otherwise stated, references to James concern his paper from 1897.

James's main example concerns religious belief. First, theoretical reasoning alone is here unable to decide the issue, or so James thinks. Moreover, given that the option is a living one, so that both alternatives are live hypotheses for the person, it is momentous and forced. It is momentous because '[w]e are supposed to gain, even now, by our belief, and lose by our non-belief, a certain vital good' (26). It is forced because '[w]e cannot escape the issue by remaining sceptical and waiting for more light, because, although we do avoid error in that way *if religion be untrue*, we lose the good, *if it be true*, just as certainly as if we positively chose to disbelieve' (26). What the religious sceptic is saying is, in effect, 'that to yield to our fear of its [the religious hypothesis] being in error is wiser and better than to yield to our hope that it may be true' (27). In other words, '[i]t is not intellect against all passions, then; it is only intellect with one passion laying down its law'. Religious scepticism represents just one passionate choice among many. It is no more rational than religious belief.

But James wants to go one step beyond merely applying this thesis which could, after all, only support the admissibility and not the unique rationality of the religious hypothesis. He proceeds to argue that the sceptical strategy is in fact less rational than the opposite alternative, since it might prevent a person from ever acknowledging certain *de facto* truths. For, assuming that religion is really true, the religious sceptic will never be in a position to acknowledge it as such, whereas a person of a more trusting nature may. On these grounds, James deems the agnostic rule for truth-seeking unacceptable. His second thesis thus states 'a rule of thinking which would absolutely prevent me from acknowledging certain kinds of truth if those kinds of truths were really there, would be an irrational rule' (28).

One could object that always to commit oneself upon insufficient evidence seems just as irrational. For, assuming that religion is not true, the trusting person will nonetheless believe it to be true. To paraphrase James: a rule of thinking which would absolutely commit me to acknowledging certain kinds of truth if those kinds of truths were *not* there, would be an irrational rule. In being agnostic we run the risk of losing the truth. In committing ourselves too hastily we run the risk of obtaining falsity. I see no reason for thinking that the one strategy be deemed objectively preferable to the other. In the

following, I shall argue that the Jamesian strategy in any case faces severe problems when applied to radical scepticism.

Did James intend his theses to be applicable also to radical scepticism? He discusses radical scepticism at several places in his will-to-believe article. What he says there strongly suggests that he would consider the issue of radical scepticism as a special case of his two theses about genuine options,⁵ and this is also how he has sometimes been interpreted in the literature.⁶ Some of James's commentators have been quite explicit on this point. Thus H. O. Mounce suggests that James's views on the matter 'are relevant wherever scepticism appears' (1997: 96). This is supposed to be true, in particular, for any scepticism that 'seeks to cast doubt on every certainty' (96), that is to say, any scepticism that has (*UC*) as a consequence. I will try to make plausible that Mounce's argument faces a dilemma that threatens to undermine any endeavour in this direction.

Before stating his argument, Mounce notes that we often claim to know things even though it is logically possible that we are in error. Suppose, to take his example, that you have parked your car somewhere, say, behind the supermarket. If you have a clear memory of where you parked it, say, twenty minutes ago, you are entitled, by normal standards, to say that you know where it is. This allows you to exclude all other possibilities as regards its location. And so when you want to go back home, you know what to do. But someone may question your knowledge claim on the grounds that, since you have not kept an eye on it since you parked it, you do not know that the police have not towed it away in the meantime. But if they have done so, it is not behind the supermarket. Consequently, you do not know where it is, and so you should suspend judgement as to its

⁵ James observes that radical scepticism cannot be refuted by logic or by theoretical reasoning alone: '[I]f a pyrrhonic sceptic asks us how we can know this [that there is a truth], can our logic find a reply? No! certainly it cannot' (1897: 10). And he claims that whether or not one is a radical sceptic is a passionate decision, i.e. a matter of volition: 'It is just one volition against another,—we are willing to go in for a life upon a trust or assumption which he [the sceptic], for his part, does not care to make' (ibid.). Moreover, he does apply his thesis to radical *moral* scepticism—to '[t]he question of having moral beliefs at all or not having them'—and goes on to say that radical intellectual and moral scepticism are on an equal footing, at least from the point of view of logic and purely theoretical reasoning: '[M]oral scepticism can no more be refuted or proved by logic than [radical] intellectual scepticism can' (ibid.).

⁶ For a recent example, see Weintraub (1997: 26). She takes James's licence to believe at our own risk any hypothesis that is live enough to tempt our will to have radical scepticism as a special case, so that, in James's view, treating everything as uncertain is just one passionate (non-intellectual) choice among many.

location. One can reason in the same way against any claim to empirical knowledge. For instance, if you seriously intend to attend a meeting this afternoon then you are entitled to say that you know you will attend it. But you do not know that you will not have a serious accident on your way to the meeting, in which case you will not be able to attend the meeting. Hence, you do not know, after all, that you will attend the meeting and you should not believe that you will.

To this Mounce responds as follows:

But in that case [if we were to treat everything as uncertain] we should be unprepared to cope even with normal circumstances, even where events take their normal course, so that had we reasoned in the normal way we should have known what to do. In practical terms, the attitude is evidently folly. If we reason in the normal way we may lose, but, on the alternative, we cannot win. Moreover, we have to adopt some attitude, since even if we evade a decision, we have in effect decided to let things drift, to treat everything as uncertain. The option, as James would put it, is forced. (100)

The choice is here presented as one of retaining all our beliefs or giving all of them up. We may grant that the option, thus presented, is forced in James's sense, and we may grant also that theoretical reason alone cannot settle the issue.⁷ It still remains to show that it is momentous and living, or James's theses will not be applicable.

There are two alternative states of the world to take into account: either events take their normal course, or they do not. Normal here means normal relative to our beliefs. Suppose that we decide to retain our normal beliefs. Then we will, as Mounce puts it, know what to do, that is to say, our beliefs and desires will suggest a certain action. If the circumstances are normal ('the car is still behind the supermarket'), my action ('go and look for it') will have a satisfactory result ('finding the car where I parked it'). If, however, circumstances are abnormal ('the police have towed the car away'), my action will lead to frustration. So if we reason in the normal way, we may

⁷ Pascal noted in his *Pensées* that the option between what he calls dogmatism (retaining our normal beliefs) and radical scepticism (giving everything up) is a forced option with respect to which one cannot remain neutral without, in effect, giving in to scepticism (paragraph 434 in the Brunschvicg edition). His final reply to radical scepticism, however, is theistic and does not take the form of a wager argument (see the same paragraph).

Table 10.1. *Consequence matrix for the sceptical wager*

	S_1 Beliefs are true	S_2 Beliefs are false
Retain beliefs	Successful action, satisfaction	Unsuccessful action, no satisfaction
Give up beliefs	No action, no satisfaction	No action, no satisfaction

‘lose’, i.e. not achieve our goals, but there is also a chance that we may ‘win’. If we decide to reason like the sceptic, by contrast, we can only lose: we will never know what to do and remain in a state of passivity.⁸ So, to paraphrase James, we can gain by our belief and lose by our non-belief a certain vital good, meaning that the option is momentous. We note that in order to reach this conclusion we made use of substantial beliefs about the consequences of acting in such and such a way under different circumstances.

Mounce does not say exactly how he would assign utility values to the different consequences, but it is not difficult to guess what the assignment would look like. The option of retaining the beliefs if they are true would be assigned a high positive number, whereas all other consequences would have more or less the same non-positive value. Here is a proposal:

$$\text{Desirability matrix: } \begin{vmatrix} 10 & -10 \\ -10 & -10 \end{vmatrix}$$

We see that we do not even have to take probabilities into account since one alternative action, that of retaining our beliefs, is a (weakly) dominating one. Retaining our beliefs is always at least as good as giving them all up.

But is this argument really coherent? To answer this question, it will be useful to observe that the sceptical option can be approached from two different perspectives. One can consider it either from our current doxastic position or from an uncommitted sceptical position.

From the point of view of our current epistemic outlook, one can indeed reason as Mounce does, at least superficially. From that point of view, the option is not only forced but also momentous. But is the

⁸ At least this is the Humean view: were we to become sceptics in real life, ‘[a]ll discourse, all action would immediately cease; and men remain in total lethargy’ (Hume 1978: 160).

option at all a live one from the point of view of our normal beliefs? What is the connection between full belief and liveliness? The fact that a given hypothesis is alive for a given person does not imply that the person in question is convinced of its truth; as James understands it, this merely means that it is 'among the mind's possibilities' (2), that there is some willingness to act (3). What we are interested in here, however, is the converse relationship. Our question is: what consequences can be drawn from the fact that a given proposition is believed for its liveliness and the liveliness of other, incompatible alternatives? James's position on this issue is that '[t]he maximum of liveliness in an hypothesis means willingness to act irrevocably', and '[p]ractically, this means belief' (3). A believed proposition, then, is maximally alive for the person who holds the belief. I take this to imply that any incompatible, belief-contravening alternative is maximally dead. That our beliefs should be false is not a living possibility from the point of view of those beliefs themselves. If I believe that my car is still behind the supermarket, it is not a live possibility for me that the police have towed it away.

James's theses are not directly applicable to the problem of radical scepticism. The situation is quite different with the religious option, which many of us naturally approach from an agnostic position, thus allowing for it to be a living option for us now. So, there is an important respect in which these two options—the sceptical and the religious—are not analogous.⁹

This raises the question whether we could not consider the sceptical option from an agnostic position instead of from our current position. This would mean, of course, to consider it from a point of view at which nothing is certain, that is, from the position of radical scepticism itself. Unlike in the religious case, this would have to be a purely hypothetical exercise. The issue now is whether we, while pretending to be sceptics, should opt for remaining sceptics also in real life or rather for sticking to our normal beliefs. This makes the option living, at least hypothetically.

⁹ James (1897: 26) writes that his argument for religious belief is directed only at those who do not rule out at the outset that the religious hypothesis is true: 'Of course, we must admit that possibility at the outset. If we are to discuss the question at all, it must involve a living option. If for any of you religion be a hypothesis that cannot, by any living possibility be true, then you need go no farther. I speak to the "saving remnant" alone.'

But the difficulty now is to argue for its momentous character. For it does not seem possible to argue, from a sceptical position, that a particular action would (normally) suggest itself to us, were we to retain our normal beliefs. There is simply nothing to appeal to in support of that claim since everything is being called into question. And nothing seems to come forth, for the same reason, regarding the consequences of acting in this or that way provided circumstances are normal. The radical sceptic denies the possibility of justifying any beliefs, including those pertaining to the causal effects of acting thus and so, or remaining inactive.¹⁰ In short, we are not in a position to claim that there is a certain vital good to be gained by our belief and lost by our non-belief. In this case, too, the Jamesian theses, be they true or not, simply do not apply to radical scepticism.¹¹

The above reasoning shows, in my view, that wager-type reasoning is unsuccessful in dealing with the sort of scepticism that bothers Bonjour and others, i.e. a scepticism that casts doubt on all our empirical beliefs. But what about other *prima facie* less dramatic sorts of scepticism, e.g. regarding natural testimony? Coady, in his book, considers and rejects the following wager argument for trusting testimony.¹² Suppose that we choose to trust testimony. If testimony is reliable but less than 100 percent so (S_1), then we will be able greatly to extend our set of true beliefs at the expense of adding now and then a false belief. If, on the other hand, testimony is unreliable but not entirely so (S_2), the result will be a great extension of error and some new true beliefs. Suppose instead that we choose not to trust testimony. Then, if testimony is reliable, we will not be in a position to take advantage of available true beliefs. If, on the other hand, testimony is unreliable, we will be greatly saved from error (see Table 10.2).

As Coady correctly remarks, 'we enter upon a very cloudy area indeed' (1992: 111) when we ask which desirability and probability

¹⁰ See Weintraub (1997: 28). Taliaferro (1992) defends a wager-type argument against radical scepticism on empirical matters. Observing that a strict sceptic cannot be assumed to know the likely outcome of his actions, he concedes that as a consequence '[a] strict pyrrhonian sceptic would be . . . at loss over how to wager' (ibid.: 229).

¹¹ There are certain parallels between James's view and Hans Reichenbach's defence of induction (1938: 358–63). For a more recent analogy, see William Alston's vindication of basic beliefs in his book from 1993. Sometimes Alston's argument takes the form of a wager (see especially p. 130). For criticism of Alston's argument, see Plantinga (2000: 117–34).

¹² The argument is presented by Coady (1992: 109) as 'an interesting reconstruction' of some elements in H. H. Price's discussion of testimony in Lecture 5 of his book from 1969, although Coady says he does not think the argument can really be found in Price's text.

Table 10.2. *Consequence matrix for the testimony wager*

	S_1 Testimony is reliable	S_2 Testimony is unreliable
Choose to believe testimony reliable	Greatly extended 'true belief' and some new false beliefs	Great extension of error and some new true beliefs
Choose not to believe testimony reliable	Great failure to get available 'true belief'	Great safety from error

Source: Coady (1992: 111).

matrices are appropriate for the testimony wager. In the absence of any better alternative, Coady suggests that S_1 and S_2 be considered 'equiprobable on the empirical evidence' (ibid.) and he goes on to propose, without argument, the following assignment of utility values:

$$\text{Desirability matrix: } \begin{vmatrix} 10 & -10 \\ -10 & 10 \end{vmatrix}$$

Coady now observes that the two alternatives will have the same expected utility given these assignments, namely 0.¹³

Although I agree with Coady's conclusion that the testimony wager does not work, I am not entirely satisfied with the way in which he reaches it. His argument hinges on specific choices of probabilities and utilities for which no detailed argument is presented. His view that we should put an equal weight on the ends of minimizing risk of error and maximizing information is only stated and not argued for. The matter is crucial since other assignments of utilities lead to other results. Suppose, for example, that one were to assign 20 and not 10 to the consequence of opting for testimony given that it is reliable in which case trust would maximize utility and hence constitute the rational alternative.

Yet, implicit in Coady's book is, I believe, another more fundamental reply to the wager which, once it is made explicit, will be seen

¹³ Standard Bayesian decision theory recommends that an action should be chosen that maximizes expected utility. The expected utility of an action is obtained by multiplying, for each possible state, the probability of that state with the utility value of the consequence of performing the action given that this state obtains, and adding the results. The expected utility of trust, for example, is calculated as follows: $0.5 \cdot 10 + 0.5 \cdot (-10) = 0$.

to mirror closely my objection to Mounce. The testimony wager, just like Mounce's sceptical wager, can be approached either from our current outlook or from a testimonially 'untainted' neutral perspective. As for the first alternative, it is part of our current outlook that testimony is considered reliable. This is not something that Coady, or as far as I know anyone else, wants to dispute. The problem is how to justify our actual trust in the word of others and not to contest its actuality. But this means that the unreliability of testimony is not a live possibility, and so there is no real choice to be made. Let us instead look at the matter from a position that is neutral with respect to the reliability of testimony. What would this neutral position look like? We recall Coady's observation, in another context, about the 'cohesion' of our informational routes and its particular repercussions for scepticism about testimony: 'Someone who sought to isolate an individualist basis in perception, memory, and inference, in order to test the reliability of testimony, would not only face the problems of language and understanding . . . but would have to discount an enormous amount of what goes into normal perception, inference, and memory' (1992: 170).

What Coady is saying is that a testimonially untainted state would be very meagre indeed, perhaps to the extent that it would contain no empirical claims whatsoever. At least, I believe a case could be made that this is a consequence of Coady's theory of testimony. But this would mean that we would lack any information that would allow us to establish with any confidence what the practical utility of trust vs. distrust will be. In particular, that we would benefit from trusting testimony, if it is generally reliable, is presumably a substantial empirical claim that does not bear much resemblance to the usual candidates for the a priori, such as 'nothing can be red and blue at the same time', 'two plus two is four', and so on. A parallel case could be made against a 'memory wager' (although I know of no such attempt in the literature). Opting for memory is either not a live option, if approached from our normal outlook, which includes a commitment to its trustworthiness, or not momentous, if viewed from a state that is neutral with respect to the reliability of memory.

Summing up this section, Jamesian wagering is ineffective against radical scepticism. Any such attempt faces a dilemma similar to that confronting Mounce's argument. Either the wager is carried out

from our everyday doxastic perspective, in which case it may be momentous but is not living, or it is approached from the point of view of radical scepticism itself, in which case it may be living but is not momentous. We have to look elsewhere for a coherent pragmatist reply to scepticism.

10.3 Peirce's Reply to Scepticism

In his criticism of Cartesian doubt, the persistent idea that we should begin in philosophy by doubting everything that admits of doubt, Peirce pointed out that it is impossible for a person to put herself in a state of universal doubt by an act of will; it is even impossible to doubt a single proposition by an act of will, if one actually believes it. To the best of my knowledge, he does not explicitly discuss the modern sceptical argument that focuses on the possibility of systematic deception. Nonetheless, as I will try to establish, it is possible to reconstruct a coherent rebuttal of that argument by combining different parts of Peirce's theory. On this reconstruction, Peirce would reject the first premiss, (S1), i.e. the premiss that we do not know that we are not deceived. In this section I will concentrate on presenting Peirce's response, as I see it.¹⁴ I will offer some additional support in its favour in the next section.

The interpretation of Peirce is a matter of scholarly dispute. Some of the interpretational issues will be addressed as we proceed. I would like to emphasize, though, that the central issue here is whether there is a coherent pragmatic response to the sceptical argument; it is not central whether Peirce, or anyone else for that matter, has actually subscribed to it.

Peirce in order to establish the falsity of (S1) needs to show that the enquirer, when he sets out in philosophy and is confronted with the sceptical argument, has no reason to doubt that he is not deceived so

¹⁴ For an alternative account of Peirce's response to scepticism, see Putnam (1995). Putnam sees the pragmatist answer to scepticism in a conception of enquiry that construes it as a 'cooperative human interaction with an environment', while rejecting the 'methodological solipsism' of Carnap and the logical empiricists (ibid.: 70). It is true that Peirce thought the ultimate problem to be 'how to fix belief, not in the individual merely, but in the society' (1877: 13). Unfortunately, Putnam's discussion is too general to bring out the supposed bearing of the social aspect of Peirce's theory on the problem of radical scepticism.

that he can hold on to his common sense belief. If this can be established, then the falsity of (S1) would follow by (S4). For (S4) entails that, if I need not stop believing that I am not deceived, then I know that I am not deceived.

First, Peirce observes that the starting point not only in daily life but also in philosophy must be the beliefs we employ in our normal enquiries and deliberations:

Philosophers of very diverse stripes propose that philosophy shall take its start from one or another state of mind in which no man, least of all a beginner in philosophy, actually is. One proposes that you shall begin by doubting everything, and says that there is only one thing that you cannot doubt, as if doubting were 'as easy as lying'. Another proposes that we should begin by observing 'the first impressions of sense', forgetting that our very percepts are the results of cognitive elaboration. But in truth, there is but one state of mind from which you can 'set out', namely, the very state of mind in which you actually find yourself at the time you do 'set out'—a state in which you are laden with an immense mass of cognition already formed, of which you cannot divest yourself if you would; and who knows whether, if you could, you would not have made all knowledge impossible to yourself. . . . Now that which you do not at all doubt, you must and do regard as infallible, absolute truth. (1905a: 167)

Second, my belief that I am a normal person with two hands and not systematically deceived is clearly part of that 'mass of cognition already formed', that is, of the beliefs I employ in my practical conduct and enquiries. As Isaac Levi, a contemporary Peircean pragmatist, puts it:

I maintain that in his practical conduct and his inquiries, X [an enquirer] is committed to dismissing the threat of a [*malin*] *genie* and to discount the worry that he is a brain in a vat. Not only is X committed to ruling out the logical possibility of wholesale deception in his firm convictions. He is committed to ruling out the logical possibility of there being even one error in his firm convictions. (1991: 58)

It follows that when I set out in philosophy, I believe that I am not systematically deceived.

Third, it still needs to be established that when I am confronted with the sceptical argument, this should not make me doubt what

I did not doubt before. Although I am committed in my practical conduct and enquiries to dismissing the threat of radical scepticism, this may be due to the fact that my belief that I am not systematically deceived has not yet been challenged. Clearly, the interesting question is what I believe, or what I am committed to believing, given that the sceptic, in the course of his argument, has denied me knowledge of non-deception and not what I believed prior to that event.

It is clear from the following passage what Peirce's response to this challenge must be:

It is important for the reader to satisfy himself that genuine doubt always has an external origin, usually from surprise; and that it is as impossible for a man to create in himself a genuine doubt by such an act of the will as would suffice to imagine the conditions of a mathematical theorem, as it would be for him to give himself a genuine surprise by a simple act of the will. (1905*b*: 484)

Peirce is here saying that only external factors may occasion doubt, and that doubt is usually preceded by surprise, implying that there may also be external doubt-inducing events that are not surprising. At another place he admits internal as well as external origins of doubt (1905*c*: 299) without explaining what those internal factors might be. Nevertheless, his most firmly held view seems to be that surprise is the only true cause of doubt, where surprise is caused by novel experience: 'For belief while it lasts, is a strong habit, and as such, forces the man to believe until some surprise breaks up the habit. The breaking of a belief can only be due to some novel experience' (1905*c*: 299).

The sceptic's denying me knowledge of non-deception is, to be sure, an external event, an event taking place in the world. But not all external events are causes for doubt. Only those that are surprising are supposed to be such causes. And no one could claim to be surprised by the fact that I am denied knowledge of the falsity of sceptical hypotheses in the course of the sceptical argument—confounded or amused, perhaps, but hardly surprised. Hence, on this analysis of what could motivate doubt, the sceptic's denying me knowledge of systematic deception should not make me doubt the validity of my knowledge claim.

10.4 More on Incoherence as a Reason for Doubt

The thesis that the sceptic's denying me knowledge of non-deception is insufficient to produce doubt is in need of more detailed argumentative backing. For one, Peirce does not provide an explicit analysis of the concept of surprise, although this concept plays a major role in his defence of the claim in question. Moreover, his thesis that surprise is the only cause of doubt seems too restrictive.

Here is a more detailed argument for the sceptic's denial of my knowledge of non-deception not being surprising, drawing on the general idea behind Paul Horwich's analysis of surprise, as outlined in Chapter 3 of the present book. What does it mean for an event to be surprising? Horwich notes, as we saw, that the fact that an event is very unlikely, given our current theory, is not sufficient to make it surprising (101). For instance, that Smith, an ordinary citizen, would win the National Lottery was antecedently very unlikely, but it is now hardly surprising that he won. After all, someone had to win, and Smith just happened to be the lucky one. What is needed in addition to low likelihood for an event to be surprising is the presence of an alternative hypothesis that would make the event likely. No credible alternative hypothesis comes to mind upon learning about Smith's newly acquired wealth. It is true that if the lottery were rigged in favour of Smith, his winning would be extremely likely. But we do not consider this a serious alternative. It does not strike us as a good alternative explanation.

Change the setting. Suppose that Smith, the winner, is in fact the brother of the Prime Minister. That the brother of the Prime Minister would win the National Lottery would strike some as surprising. Why? Because the explanation that the lottery has been rigged in favour of the brother of the Prime Minister may occur to some as a viable alternative to the assumption of fairness.

This is, by the way, where Bonjour went wrong when he proposed 'probabilistic inconsistency' as something that always leads to incoherence. Contrary to what his second coherence criterion states, probabilistic inconsistency is not by itself incoherence-inducing. Incoherence results only if there is an alternative explanation to the seemingly unlikely event, an explanation that would render that event likely.

Let us apply this to the sceptical scenario. For me to be surprised so as to give up my belief that I am not systematically deceived there must be some novel experience that is very unlikely on that hypothesis but at least to some extent likely on some deception hypothesis. Now we usually would not give any credit whatsoever to radical deception hypotheses. But even if we did, no deception hypothesis could make my experience more likely, for deception hypotheses are all, *ex hypothesi*, phenomenologically indistinguishable from everyday life, and so whatever is unlikely given my belief that I am an actual person is equally novel given, say, the alternative hypothesis that I am a brain in a vat. Hence, I will never doubt that I am not systematically deceived *as the effect of being surprised*. If surprise is the only cause of doubt, as Peirce sometimes suggests, then nothing could make me doubt that I am not deceived.

Of course, this goes in particular for the ‘novel experience’ deriving from the sceptic’s denying me knowledge of non-deception. The event of the sceptic’s asserting that I am in fact a mere brain in a vat is no more likely on the brain-in-the-vat hypothesis than it is on the usual common sense picture.

Still, Peirce’s thesis that surprise is the only cause for doubt is not quite correct. Isaac Levi has noticed that we sometimes have reasons to doubt a theory because it fails to explain events which we want it to explain (1991: 153). An anomaly, while normally being prompted by external circumstances (by the observation of events not explained by the theory), does not necessarily involve an element of surprise. For, the unexplained events involved in the anomaly, unlike surprising events, need not be unlikely given the theory. While not all anomalies are caused by surprising events, all surprising events are arguably anomalous.

This notwithstanding, explanatory failure is, as Levi also points out, not incentive enough to give up the current view: ‘One should not question a theory concerning whose truth one is currently certain merely because it has not yet proved capable of explaining . . . some phenomenon’ (ibid.). The reason is that to give up a theory deprives one of the information contained in the theory that normally has some value in explaining other phenomena (ibid.). Rather, we need the additional impetus of an alternative candidate theory (ibid.: 154). The new theory, moreover, must have some promise as a means for removing anomaly while still being capable of explaining everything that could be

explained by the old theory (*ibid.*). Indeed, the alternative theory must represent a genuine improvement relative to the current theory: ‘To justify opening up one’s mind to the hypothesis *A* initially taken not to be seriously possible, the result of ending up with *A* in the corpus must be more valuable informationally than status quo’ (*ibid.*: 157).

If surprise is the only cause for doubt, then I do not doubt that I am not deceived, so I believe that I am not. But surprise, as we just saw, is not the only reason for doubt. Anomaly can also have this effect. So, does the sceptic’s denying me knowledge of non-deception produce an anomaly in my view, and if so, does this present a reason for doubt? The answer is this: sceptical hypotheses, by definition, share all observational consequences with our normal view. Hence, they will not be able to explain events that could not be explained before. This means that they are unable to remove anomaly from the current theory. In other words, no anomaly in our view, be it due to the sceptic’s denying our knowledge of non-deception (if this indeed induces anomaly) or to other events, could ever make us lend sceptical hypotheses a charitable ear.

While Levi should be credited for bringing in the notion of anomaly, his own short discussion of scepticism is incomplete on two accounts. First, he never raises the issue of what happens in the particular circumstance where the sceptic has denied me knowledge of non-deception. Nevertheless, as I have argued above, it follows from his theory of ‘uncoerced contraction’ that it would be illegitimate to contract in this case so as to give the sceptic a neutral hearing, since no anomaly or surprise is induced by his assertion (1991: section 4.9). This point is far too important to be left implicit in Levi’s work. Second, to the best of my knowledge Levi has not provided an explanation of why many philosophers are prone to scepticism.

Anomaly in Levi’s sense means ‘explanatory gap’. Something we would like to have explained cannot be explained by the current theory. This may, as I have argued, in some circumstances occasion doubt. There is also another sort of anomaly—lacking a better term let us call it incongruence—that can provide a reason for doubt. Consider the following example:¹⁵ suppose I believe that *X* stole my money. I believe this because only *X* had a key to my apartment, and I know he was in financial difficulties. Then my neighbour reminds me that *Y*’s balcony adjoins mine (which would give him

¹⁵ I owe the example to an anonymous referee.

equally good access to my apartment) and that *Y* needed money, too. These facts need not surprise me. I may have known them all along, but just not thought hard enough about the matter. And there are arguably no anomalies in Levi's sense here. The problem, vaguely speaking, is rather that one part of my belief system (the one about *X*'s key and his financial problems) points in one direction (toward *X*'s and against *Y*'s guilt) and another part (the one about *Y*'s balcony and his financial problems) in another direction (toward *Y*'s and against *X*'s guilt). It would be reasonable in these circumstances to start doubting the hypothesis about *X*'s guilt so as to give the hypothesis about *Y*'s guilt a neutral hearing.

Allowing for such incongruence to be doubt-inducing does not mean letting the sceptic in through the backdoor. That is to say, it does not mean that a person should start doubting his common sense views when presented with a sceptical hypothesis. The sceptical predicament is structurally different from the criminal story just told. For the sceptical case to be structurally similar two conditions would have to be satisfied: (i) one part of a person's beliefs would have to support the common sense view while counting against, say, the brain-in-the-vat hypothesis; and (ii) another distinct part would have to support the brain-in-a-vat hypothesis while counting against our common sense picture. Clearly (i) is trivially true: the part of a person's beliefs containing his common sense convictions uniquely supports the common sense picture of the world. By contrast, (ii) is surely false. There is, I submit, no part of a person's belief system that supports the brain-in-a-vat hypothesis and at the same time disconfirms the common sense view (what would such a system look like?). I am inclined therefore to think that sceptical hypotheses do not present us with the kind of incongruence that we are concerned with here.

In the next section I will sketch an account of how scepticism emerges, but first I will raise the question of what is distinctively pragmatic in the Peirce-inspired theory.

10.5 Three Roads to Scepticism

What is 'pragmatic' about the Peircean approach? First, Peirce insists that we have to set out in philosophy from our normal daily-life perspective. The function of our normal beliefs is to provide what

Levi calls a standard of serious possibility, the overarching purpose of which is, arguably, to guide action so as to satisfy our desires. The practical perspective is, in this sense, primary. Second, the conception of knowledge underlying Peirce's approach is practical in the sense of being 'purpose-driven'. For it sees knowledge as the goal of enquiry, which is a purpose-driven activity, rather than as only vaguely related to human purposes, aims, and practices.

Third, the criteria we use when we decide whether or not to give an alternative hypothesis a hearing are ultimately practical in nature. If all we cared about were avoiding error, we would be better off by becoming sceptics. But avoiding error is not the sole goal of enquiry; we also want to arrive at substantial truth. James was quite clear on the last point. In an early paper from 1882, 'The Sentiment of Rationality', he points to our wish to 'banish uncertainty from the future' (77) so as to 'define expectancy' (81). The objective to define expectancy, moreover, he classified as ultimately practical because of its significance in practical action. John Dewey similarly emphasized our 'quest for certainty', adding that 'the ultimate ground of the quest for cognitive certainty is the need for security in the results of action' (1929: 39). Unless we have a clear sign that our current view contains error, we are not willing to sacrifice its predictive force.

How can we explain, on the pragmatist model, that some philosophers claim they do not know they are not brains in vats (or otherwise deceived)? A plausible analysis of scepticism should be able to account for this fact. On the Peircean model, the attraction of scepticism can be explained by reference to a failure to observe one or more of the following three characteristics of rational enquiry.

In enquiry the starting point is our ordinary beliefs rather than, say, a *tabula rasa*. It is true that from the standpoint of a *tabula rasa* the normal view and any given sceptical hypothesis are symmetrical: they have the same observational consequences and there is no reason to prefer the one to the other. There is no escape from scepticism from the standpoint of a *tabula rasa*. But when we approach the argument of the sceptic, we do so from our ordinary doxastic perspective, in the absence of any good reason to do otherwise. From that point of view, the normal view and the sceptical hypothesis are not symmetrical: the former is believed and the latter is not.

However, acknowledging the primacy of our ordinary perspective is no vaccine against scepticism. In philosophy we are sometimes too quick to consider alternative theories on a neutral basis. Indeed, we usually do so automatically and, as a consequence, we tend to forget that there is at all a stage in enquiry preceding the neutral examination of alternative views, a stage at which it is decided whether or not to give an alternative theory a hearing in the first place. The crucial point is that this decision is made on the basis of an evaluation carried out from the current doxastic position and not from a neutral perspective. Scepticism is the inevitable result of jumping directly from our daily life perspective to a neutral examination of radical doubt, without first contemplating the possible gain in so doing. In particular, the sceptic's denying me knowledge of non-deception is insufficient to occasion genuine doubt.

Finally, scepticism exerts a particularly strong attraction on philosophers who value only avoidance of error at the expense of the acquisition of substantial information. From the point of view of error avoidance, radical doubt is certainly optimal: the best way to avoid error is to believe nothing at all. However, this position is purely hypothetical, as every person with desires and purposes must value attaining substantial information which can be used to guide her actions so as to satisfy those desires. There is, nonetheless, a tendency in philosophy to assume that the practical perspective is unimportant.¹⁶ This suggests a third form of negligence: the refusal to take our interest in practical action into account when deciding whether or not to open up for a new hypothesis. This, too, is a straight road to scepticism.

It should be mentioned that what has been said so far about scepticism regarding the beliefs themselves applies equally to scepticism regarding one or more of the routes to belief, such as testimony or observation. As Peirce noticed, the point of departure in philosophy is our normal cognitive state which is a state laden with an 'immense mass of cognition already formed'. Clearly this state includes not only beliefs already formed but also informational routes already trusted, as exemplified by our reliance on memory and the word of others. What are we to do, then, with a sceptic who invites us to debate

¹⁶ Cf. Mounce (1997: 101).

the reliability of one or more of these routes on 'neutral ground'? One thing is for sure: the sceptic's denying us knowledge of their reliability is by itself insufficient to motivate neutral enquiry. We may take on the sceptic's challenge only if we see some prospects of a gain in so doing. But it is difficult to see how we could possibly benefit from this undertaking. In the worst case we lose not only a long trusted source of reliable belief but also many of the true beliefs that we have derived from it in the past. The most we can hope for, if we succeed somehow in justifying our trust in a way that would satisfy the sceptic, is to regain what we already had before we accepted the sceptic's invitation. The appropriate response would seem to be not to accept the sceptic's challenge after all. This reasoning in terms of consequences and utilities is perfectly legitimate in this case since, as I have argued, whether or not to engage in neutral examination of our present commitments is a decision we make from the current doxastic standpoint and not from a neutral perspective, and from the standpoint of the former all actually held beliefs are true and all actually trusted sources trustworthy. We have to decide whether to accept the offer on the basis of how things seem to us now, i.e. how they appear to us from our current cognitive state.¹⁷

When I present the Peircean ideas to mainstream epistemologists I invariably encounter strong resistance or even indignation. Before I proceed I would like to take the opportunity to reply to some common objections that are based on what I take to be misunderstandings.

Objection: The present proposal is absurd. According to it, all I have to do is stubbornly hold on to my ordinary beliefs, and by dint of that alone, I can 'refute' the sceptic, just by telling him that I still hold these beliefs. If this is acceptable, why could I not respond to any challenge to any belief in this manner, provided that I can be sufficiently dogmatic to just keep holding on to my present beliefs? For instance, theists might respond to the argument from evil by saying: 'Well, I have not stopped believing in God after hearing your argument. Therefore, I know there is a God.'¹⁸

¹⁷ This outline of the Peircean strategy is certainly incomplete in many respects. Isaac Levi has addressed some of the remaining issues in his book from 1991.

¹⁸ The objection was raised by an anonymous referee.

Reply: Far from advocating dogmatism, I have, following Peirce, pointed to several legitimate reasons for doubt: surprise, anomaly, and incongruence. In these cases doubt is not only legitimate but, arguably, obligatory. A person who sticks to his old beliefs in spite of severe intellectual dissonance could be criticized for so doing. We are accordingly not always free to hold on to our present beliefs. As for the theist example, the issue is not whether I have actually stopped believing in God after hearing the argument from evil. It is rather whether I *should* stop. That in turn depends, again, on whether the argument in question succeeds in inducing the required sort of cognitive dissonance.

Objection: The claims about what are the only possible ways of coming to doubt something seem speculative and themselves doubtful. For example, it seems that it would be possible to come to doubt p by being presented with an argument to the effect that the grounds one had for believing p were not really adequate.¹⁹

Reply: I believe that many cases of being presented with an argument to the effect that one's grounds were not really adequate can be subsumed under the different forms of cognitive dissonance that have been discussed above: surprise, anomaly, and incongruence. Suppose for example that I believe that X stole my money because I believe the following: (i) X had the opportunity, (ii) X belongs to a certain ethnic group, and (iii) members of that group are particularly prone to stealing. As I turn on the TV, an internationally acclaimed expert on criminal statistics flatly denies that there should be any such correlation at all, citing empirical evidence in support of her denial. This would be a case of reliable questioning of one of my premisses. I agree that in this case it would be reasonable for me to cease believing that X stole my money. But this case can be accommodated within the present theory. For surely hearing the expert deny my firm belief is a surprising event for me. It is surprising because, first, given my firm belief that it is true that members of the group in question are prone to theft it was initially very unlikely that a reliable expert would deny the claim. After all, being reliable means being disposed to tell the truth. Second, there is an alternative hypothesis available, one that would explain why the expert rejected the

¹⁹ I owe the objection to an anonymous referee.

supposed correlation. The alternative hypothesis is simply that there is no such correlation so that her claim is true.

Objection: The worry is that the Peircean response discourages much of what we value in academic pursuits. Academics tend to value freedom to go down some unexplored route, do some ‘out-of-the-box thinking’, although there are no immediate benefits to be seen. We believe that good things happen when we dare to go off the tracks of our everyday doxastic positions, give alternative hypotheses a hearing (even if it is not quite clear yet why they should get a neutral hearing), and we think that we should be driven not by practical concerns, but by a desire for knowledge for its own sake. So the Peircean response to scepticism is one that goes against the self-proclaimed ideals of academic enquiry.

Reply: Perhaps the objection has its roots in a conflation of real doubt with mere hypothetical doubt. Nothing of what I have said has any bearing on the legitimacy of engaging in various sorts of hypothetical reasoning and thought-experiments. In particular, I have not said anything about hypothetical doubt. The sort of doubt that I am concerned with, following Peirce, is the real thing. The thesis is that real doubt comes from some sort of cognitive dissonance (surprise, anomaly, or incongruence). We may of course imagine that we doubt this or that—for no special reason at all. That is fine with me and, I believe, with Peirce. Perhaps I can even imagine myself doubting all my common sense beliefs at once. Still, that would not mean that I *actually* doubt all my common sense beliefs. The bottom line is that the Peircean theory does not in any way curb the imagination of scientists. At the same time it does justice to another side of scientific practice: the tenacity with which scientists hold on to their current theory. Scientists usually do not give up their current doctrine as soon as an alternative view is on the table. They do however often start to consider alternatives seriously in the face of cognitive dissonance occasioned by surprising observation, anomaly, or incongruence.

10.6 Comparison with Other Contemporary Responses

The next undertaking will be to compare the pragmatic response to scepticism to the currently most popular approaches. I do not pretend

to give a full coverage of this issue here and some of my remarks will be rather superficial. I believe that a rough comparison with other responses can still serve the purpose reasonably well of highlighting the distinctive features of the Peircean approach.

Let us return to the original sceptical argument.

(S₁) I do not know not-*SH*.

(S₂) If I do not know not-*SH*, then I do not know *O*.

Hence:

(S₃) I do not know *O*.

The pragmatist avoids the sceptical conclusion by rejecting (S₁). We do know that we are not brains in vats or systematically deceived. This proposal, as we have seen, can be supported by a principled account of knowledge and enquiry.

Let us first consider the prospects of rejecting (S₂). Philosophers of this inclination argue that, despite appearances, it does not follow from the alleged fact that we lack knowledge of the denials of radical sceptical hypotheses that we thereby lack knowledge of ordinary propositions as well. One approach is the so-called relevant alternative line of argument, according to which sceptical error-possibilities are just not relevant to everyday knowledge in the way that everyday error-possibilities are. The problem is to support this contention with a principled account of why sceptical error-possibilities are irrelevant. Fred Dretske (1970) has made important contributions to this aspect of the theory. Recent discussion has focused on the so-called Closure Principle.

(Closure Principle) If *S* knows that *p*, and *S* knows that *p*-entails-*q*, then *S* knows that *q*.

For instance, if one knows the ordinary proposition that one is currently seated, and one further knows that if one is currently seated then one is not a brain in a vat, then one must also know that one is not a brain in a vat. Several attempts have been made to construct theories that violate closure, one prominent example being Nozick's causal theory of knowledge (Nozick 1981).

Yet there is a *prima facie* tension involved in adopting such a proposal since the intuition that we know the known logical consequences of what we know is 'extremely strong' (Pritchard 2002: 222).

Pragmatism, as I construe it, does not require that we reject the Closure Principle. It leaves the sceptic's second premiss, (S₂), untouched.

Contextualism is another response to radical scepticism that has attracted recent attention. Its point of departure is our presumed 'biperspectivalism' (Williams 1991): on the one hand, we tend to think that in everyday contexts it seems perfectly appropriate to ascribe knowledge to subjects; on the other hand, ascription of knowledge to subjects in conversational contexts in which sceptical error-possibilities have been raised seems improper. While scepticism seems compelling under the conditions of philosophical reflection it is never able to affect our everyday life where it is all but ignored. How can this be explained?

The contextualist suggests that these intuitions are not conflicting but rather the effect of a responsiveness on the part of the knowledge-attributor to a fluctuation in epistemic standards caused by a change in the conversational context. More specifically, the strength of an epistemic position that an agent needs to be in if she is to have knowledge can vary from context to context. This, if true, would allow for the possibility that it can be true both that one has knowledge in everyday contexts and that one lacks it in sceptical conversational contexts. The most influential advocates of this sort of theory are Steward Cohen (1986), Keith DeRose (1992; 1995), and David Lewis (1979; 1996).

The contextualist needs an independent account of what can cause a change in conversational context, and this account must imply that the sceptic's denying someone knowledge of non-deception is sufficient to change the context into a more demanding one. A detailed theory designed to accomplish this has been presented by DeRose, who characterizes the mechanisms that raise our epistemic standards as follows: 'When it is asserted that some subject *S* knows (or does not know) some proposition *P*, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to require *S*'s belief in that particular *P* to be sensitive for it to count as knowledge' (1995: 36). *S*'s belief in *P* is sensitive if *S* does not believe *P* in the nearest possible world in which *P* is false. So, for instance, if someone asserts that I do not know that I am not a brain in a vat, this raises the standards for

knowing to such a level as to require my belief in 'I am not a brain in a vat' to be sensitive to count as knowledge, so that in the nearest possible world in which I am a brain in a vat, I do not believe that I am not a brain in a vat. Since, in the nearest brain-in-a-vat world, I presumably still believe that I am not a brain in a vat, I do not know that I am not a brain in a vat in this new context.

There are some well-known problems with the contextualist proposal. Perhaps the most serious difficulty is that, at least in the version advocated by David Lewis, it legitimizes the concern that the sceptic's standards are the right standards and that we lack knowledge after all. Lewis (1979) argues that 'knowledge' is context-sensitive in the same way in which terms like 'flat' are. We may all agree that the table in front of us is 'flat' in an everyday context. But if someone enters the room and denies that it is flat we do not thereby disagree with her. Instead, we take it that she means 'flat' in some more demanding sense and so raise the standards for 'flatness' so as to make her assertion true. Lewis calls this mechanism a 'rule of accommodation'. Now there is a sense in which the more demanding standard of 'flatness' is the more correct, scientific one, whereas our everyday assessment is just loose talk. If knowledge is construed along the same lines, the sceptic comes out as using a more precise and correct standard for assessing knowledge claims.²⁰

Pragmatism, as I understand it, rejects the biperspectivalism that was the starting point for the contextualist, thereby undercutting the motivation for thinking that knowledge is context-dependent in the first place.²¹ Maybe ascription of knowledge to subjects in conversational contexts in which sceptical error-possibilities have been raised seems improper. Still, this is mere appearance and can be

²⁰ For this point, see Pritchard (2002: 229). Williams has developed a version of contextualism that does not concede to the sceptic that she employs a higher epistemic standard.

²¹ Mounce (1997: 101) disagrees: 'Further reflection will now reveal that there is no incompatibility between the perspective adopted by the sceptic and that of ordinary practice, *just so long as neither is taken as absolute*. It is not the perspective that the sceptic adopts but his refusal to accept any other that occasions his scepticism. There is no compatibility in admitting that perspective along with that of ordinary practice. Indeed it can be salutary to do so. It is sometimes useful to remind oneself that our normal expectations are not guaranteed, that our certainty is not absolute.' For a compelling criticism of the use of such 'double standards', see Levi (1991: 58–60). Levi's conclusion is that 'the double standard either collapses back into a single standard or ensures the irrelevance of the second [sceptical] standard for the purposes of practical deliberation or theoretical enquiry' (ibid.: 60), and so 'to say that it, nevertheless, has philosophical importance is to cast doubt on how important philosophical importance can be' (ibid.). One could add that if pragmatism is construed as using double standards, it loses its distinctive character and becomes difficult to separate from contextualism.

explained by reference to the three types of negligence described above, i.e. (i) the failure to appreciate that our everyday doxastic position is the starting point of all enquiry, or (ii) the tendency to skip, perhaps unconsciously, the step in enquiry at which it is decided whether an alternative hypothesis should be given a neutral hearing, or (iii) the disregard of our legitimate interest in substantial truth at that step.

For example, the mere fact that we have been denied knowledge of non-deception is not reason enough to make us doubt that we are not deceived. Failure to see this may lead one to think that ascription of knowledge to subjects in contexts where sceptical hypotheses have been raised is inappropriate. For us to open up for an alternative hypothesis, the current view has to be felt to be unsatisfactory, perhaps due to the presence of an anomaly, and there has to be some hope that the alternative hypothesis could improve upon the situation. As I have tried to show, accepting a sceptical hypothesis never does represent a genuine improvement in comparison to the normal, non-sceptical view.

The pragmatist approach bears a superficial similarity to G. E. Moore's refutation of scepticism. Moore argued, on the basis of his conviction that the conclusion of the sceptical argument must be false, that he knows the denial of the sceptical hypothesis in question (Moore 1925). However, he failed to supplement this strategy with an analysis of knowledge that would support it. So-called neo-Moorean theories attempt to provide an account of knowledge that would establish the falsity of (S1). A general problem for neo-Moorean theories, though, is their 'lack of diagnostic appeal' (Pritchard 2002: 239), that is to say, they have trouble explaining why so many philosophers have felt that radical scepticism is somehow intellectually unavoidable.²² For the pragmatist this is no mystery. From her standpoint, there are three roads to scepticism corresponding

²² For an exception see Sosa (1999). Sosa suggests that we regard what he calls 'safety' as necessary for knowledge. This notion is characterized as follows: an agent's knowledge of a contingent proposition, *A*, is safe if the agent has a true belief in *A* in the actual world and, across a wide range of nearby possible worlds, if she believes *A*, then *A* is true. Sosa argues that the requirement of safety has some important advantages over sensitivity in DeRose's sense. Without going into any details, Sosa's proposal allows one to reject (S1) in a principled manner. Now Sosa argues that 'I don't know not-SH' might be appealing, though false, because the safety requirement on knowledge is all too easily confused with sensitivity, and we have already seen that I do not know not-SH if sensitivity is imposed as a requirement on knowledge.

to neglecting one of the three pragmatic aspects of enquiry noted above.

10.7 Conclusion

It is time to sum up the discussion in this chapter. I have tried to shed some light on the pragmatist response to radical doubt, focusing on two responses to be found in the American pragmatist tradition differing not only in the details but in the very argumentative strategy employed. One is James's attempt to turn his argument against scepticism in religious matters into an argument against radical scepticism. This proposal, however, did not survive careful scrutiny. Peirce's response, on the other hand, was seen to contain the core of a tenable pragmatist response to radical doubt.

In its essence, the Peircean response amounts to this: the sceptic insists that I should justify my belief that I am not a brain in a vat without begging the question against someone claiming that I am. If I cannot do so, the sceptic claims that I do not know that I am not a brain in a vat. The pragmatist response is that since I am convinced that I am not a brain in a vat, and, unless I have good reasons to open up my mind to his view (that is to say, to come to doubt my own view), there is no rational basis for me to consider it. The sceptic cannot offer his mere disagreement as a sufficient warrant for my opening up my mind. The onus is on the sceptic to give me justification for considering his challenge.

The Peirce discussion was particularly important for our purposes not only because it suggests a plausible reply to scepticism but also because it highlights the role of incoherence in our enquiries. Coherence may not suffice to justify our beliefs, but incoherence is what forces us to give them up.

It has been suggested to me that the characteristic pragmatist response to radical scepticism is 'to dismiss it at the outset'. If this means to disregard it without philosophical argument, it is unlikely to convince anyone. It is also unnecessary. As I have tried to make plausible, there is a coherent and distinctively pragmatic response to radical scepticism. It is as precise and detailed as the better known rebuttals—such as contextualism and relative alternatives

theory—without collapsing into any one of them. Indeed, I think I have provided some reasons for thinking that the pragmatist proposal is, in some respects at least, even superior to those other theories in the form in which they have been defended by some of their most influential proponents, although there is certainly more to be said about this matter. Saying it will have to await another occasion.

Appendix A

Counter-example to the Doxastic Extension Principle

Let \mathbf{S} and \mathbf{S}' be two testimonial systems. \mathbf{S}' is a non-trivial extension of \mathbf{S} if and only if (i) $\mathbf{S} \subset \mathbf{S}'$ and (ii) the content of \mathbf{S}' is logically stronger than the content of \mathbf{S} (i.e. there are logical consequences of the contents of \mathbf{S}' that are not logical consequences of the contents of \mathbf{S}). The Testimonial Extension Principle says this:

(TEP) If \mathbf{S}' is a non-trivial extension of \mathbf{S} , then \mathbf{S}' is less probable than \mathbf{S} .

As a special case, we have the Doxastic Extension Principle:

(DEP) If \mathbf{S} and \mathbf{S}' are doxastic systems and \mathbf{S}' is a non-trivial extension of \mathbf{S} , then \mathbf{S}' is less probable than \mathbf{S} .

As I have argued, Klein and Warfield need (DEP) to be true for their counter-example to the truth conduciveness of coherence to work. I will now give a counter-example to this principle. The example shows that a non-trivial extended doxastic system can in fact be more probable than the original system.

Let us consider the following simple variation on the Dunit example. Suppose that there has been a robbery. A conscientious detective would like to know whether Dunit committed the robbery (R) and consults independent witnesses in order to gather evidence. Although the witnesses need not be fully reliable, they all have a track record of being sufficiently reliable, and our detective routinely adopts the belief that some item of evidence holds just in case there is a witness report to this effect.

As was noted in the very beginning of this book, we often adopt beliefs in a routine-like manner. As Isaac Levi (1991: 71) describes the process, '[i]n routine expansion, the inquirer expands according to a program for adding new information to his state of full belief or corpus in response to external stimulation'. A characteristic feature of such routine expansion is that it does not rely on inference. While routine expansion does begin with certain

assumptions or premisses—namely, assumptions about the reliability of the programme—the expansion adopted is not inferred from these premisses (ibid.: 74). On Levi’s view, routine expansion includes consulting witnesses (ibid.: 75). As we have seen, BonJour emphasizes the importance of such automatically acquired (‘cognitively spontaneous’) beliefs for a coherence theory.

We will assume that each item of evidence is reported by one single witness. After querying a bystander, the detective adopts the belief that Dunit was driving his car away from the crime scene at high speed (C). After querying one of Dunit’s neighbours, he adopts the belief that Dunit is in the possession of a gun of the same type as the one used in the robbery (G). The original doxastic system \mathbf{S} contains the pairs $\langle BelC, C \rangle$ and $\langle BelG, G \rangle$. Subsequently, a new witness steps forward: after querying the bank clerk in Dunit’s bank, the detective adopts the belief that Dunit deposited a large sum of money in his bank the day after the robbery (M). Dunit’s non-trivially extended doxastic system \mathbf{S}' now contains the additional pair $\langle BelM, M \rangle$.

One might object that we could have discussed an extension from one to two beliefs, rather than from two to three beliefs. From a mathematical point of view, this would indeed have been sufficient for the purpose of rejecting (*DEP*). However, from an epistemological point of view, the one-proposition case is, as we noted already in Chapter 2, problematic. The reason is that we want to compare two belief systems with respect to their coherence, and one belief hardly qualifies as a belief *system*. More fundamentally, coherence is, as we also saw, a concept that simply does not apply to singletons. This is what Rescher’s Principle says. Hence, we cannot compare a singleton with a set of two or more propositions with respect to coherence. The relation ‘more coherent than’ is undefined if one of the relata is a singleton. In order to avoid such conceptual problems, we consider an extension of a belief system from two to three propositions.

It is easy to construct a case in which (*DEP*) is false on the grounds that the extended doxastic system \mathbf{S}' is *equally* probable as the original doxastic system \mathbf{S} . Suppose that the witnesses are all *fully* reliable. Then, obviously,

$$(*)P(C, G/BelC, BelG) = 1 = P(C, G, M/BelC, BelG, BelM)$$

Information that derives from fully reliable informants is always maximally likely to be true. Hence, if all informants are maximally reliable, then the extended system will be as probable as the original. This observation alone is sufficient to disprove the (*DEP*) as it stands.

But (*DEP*) is actually stronger than it needs to be. What Klein and Warfield need is only a weaker principle to the effect that a non-trivially extended doxastic system is not more probable than the original. Let us refer to this as the Weak Doxastic Extension Principle (*WDEP* Kern). That is the principle we need to rebut. The challenge, then, is to construct a case in which the extended doxastic system \mathbf{S}' is *more* probable than the original doxastic system \mathbf{S} .

First, let us assume that there is a large number n of suspects and that each suspect stands an equal chance of having committed the robbery, so that

$$(i) P(R) = 1/n = 1 - u, \text{ for } u \approx 1, \text{ but } u \neq 1$$

Second, let us assume that, although the witnesses are highly reliable, they are less than fully reliable: there is a small chance that a bystander report is forthcoming to the effect that Dunit was speeding away from the crime scene, although he actually was not, and there is a small chance that no bystander report is forthcoming to the effect that Dunit was speeding away from the crime scene, although he actually was. Similarly, for the two other witnesses (in the following I use P as a variable for propositions and p, q, s, t , and u as numerical variables):

$$(ii) P(\text{Bel}P/P) = p \text{ and } P(\text{Bel}P/\text{not-}P) \\ = 1 - q \text{ for } p, q \approx 1, \text{ but } p, q \neq 1 \text{ and } P = C, G, M$$

Third, we permit probability distributions that leave a small chance that the evidence is misleading. There might be a small chance that Dunit committed the robbery and slipped away on the subway or that Dunit did not commit the robbery, but just happened to be speeding at the wrong time at the wrong place. Similarly for the other items of evidence:

$$(iii) P(P/R) = s \text{ and } P(P|\text{not-}R) \\ = 1 - t \text{ for } s, t \approx 1 \text{ and } P = C, G, M$$

The next step is to introduce some assumptions of probabilistic independence. These assumptions are introduced here mainly to simplify calculations, but it is interesting to note that they characterize a common type of information-gathering involving independent evidence, independent witnesses, and an obliquely testable hypothesis. What this means is explained below.

Independent Evidence. The respective items of evidence are probabilistically independent of any other items of evidence, conditional on the

hypothesis. What does this mean? Suppose that we actually know whether Dunnit committed the robbery. Then there is a certain chance that he was speeding on the motorway away from the crime scene. Now suppose that we learn in addition that he is in possession of a gun of the same type as the one that was used in the robbery. Then learning this new item of evidence will not affect the chance that Dunnit was speeding on the motorway. Note that this assumption is not always fulfilled: the items of evidence may be of a nature that does not warrant this assumption. For instance, the fact that Dunnit came into the repair shop the day before the robbery for a tune-up so that his car would perform optimally at high speed would also constitute an item of evidence that he committed the robbery, but it would not constitute an independent item of evidence.

Independent Witnesses. The detective's routinely acquired belief about some item of evidence is probabilistically independent of any other item of evidence or of any other of his routinely acquired beliefs, conditional on the evidence. What does this mean? Suppose that we actually know whether Dunnit was speeding on the highway. Then there is a certain chance that a reliable witness would step forward with a report to this effect and that the detective would adopt this report as a routinely acquired belief. Now suppose that we learn in addition that Dunnit was in possession of a gun of the same type as the one that was used in the robbery or that there was a witness report to this effect. Then this will not affect the chance that a reliable witness would step forward with a report to the effect that Dunnit was speeding. This assumption stipulates that each witness is focused on the items of evidence that he reports on and does not attend to other items of evidence or to reports about other items of evidence. What would it take for this condition not to be fulfilled? Suppose that the witnesses have been doing their own detective work: they checked out other items of evidence or talked to witnesses who reported on these items of evidence. Then their judgements of whether it was really Dunnit they saw speeding in the car could well be affected by whether there was other evidence to the effect that Dunnit committed the crime or by whether there were any reports of such evidence.

Obliquely Testable Hypothesis. The detective's routinely acquired beliefs about the evidence are probabilistically independent of the hypothesis, conditional on the evidence. What this means is that none of the witnesses has any direct access to whether Dunnit committed the robbery or not. This question remains hidden in the black box: the witnesses' only access to it is through the items of evidence. Suppose that we actually know that the items of evidence obtain (or that some or none obtain). Then there is a

certain chance that reliable witnesses will step forward and that the detective would come to acquire beliefs to the effect that the evidence obtains. Now suppose that we learn in addition that Dunit actually committed the crime. Then this will not affect the chance that the detective would come to acquire beliefs to the effect that the evidence obtains. What would it take for this condition not to be fulfilled? Suppose that the witnesses got a quick glimpse of the robbery scene. Then their judgements of whether it was really Dunit who was speeding on the motorway, of whether Dunit really has the same gun as the one used in the crime scene, and so on, may be coloured by what they were able to gather from the robbery scene.

These independence assumptions can be expressed formally in the style of Dawid (1979) and Spohn (1980). For the technical details, I refer to Bovens and Olsson (2002).

We can now show that the Weak Doxastic Extension Principle is false for this particular example for some plausible values of p , q , s , t , and u . The probability of the original doxastic system \mathbf{S} is:

$$(\text{Prob}_{\mathbf{s}}) P(C, G / \text{Bel}C, \text{Bel}G)$$

The probability of the extended belief system is:

$$(\text{Prob}_{\mathbf{s}'}) P(C, G, M / \text{Bel}C, \text{Bel}G, \text{Bel}M)$$

We apply Bayes's theorem to $(\text{Prob}_{\mathbf{s}})$:

$$(\text{Prob}_{\mathbf{s}.1}) \frac{P(\text{Bel}C, \text{Bel}G / C, G) P(C, G)}{P(\text{Bel}C, \text{Bel}G)}$$

which equals

$$(\text{Prob}_{\mathbf{s}.2}) \frac{P(\text{Bel}C, \text{Bel}G / C, G) \sum_{\mathbf{R}} P(C, G, \mathbf{R})}{\sum_{C, G, \mathbf{R}} P(\text{Bel}C, \text{Bel}G, C, G, \mathbf{R})}$$

I am here using bold italicized letters as random variables in the statistical sense. For instance, \mathbf{R} can take on either of the two possible truth-values of R , and so on. We apply the chain rule:

$$(\text{Prob}_{\mathbf{s}.3}) \frac{P(\text{Bel}C, \text{Bel}G / C, G) \sum_{\mathbf{R}} P(C, G / \mathbf{R}) P(\mathbf{R})}{\sum_{C, G, \mathbf{R}} P(\text{Bel}C, \text{Bel}G / C, G, \mathbf{R}) P(C, G / \mathbf{R}) P(\mathbf{R})}$$

Our assumptions of conditional independence now permit us to to simplify (for the details, see Bovens and Olsson 2002):

$$(Prob_{s.4}) \frac{P(BelC/C)P(BelG/G) \sum_{\mathbf{R}} P(C/\mathbf{R})P(G/\mathbf{R})P(\mathbf{R})}{\sum_{\mathbf{C},\mathbf{G},\mathbf{R}} P(BelC/\mathbf{C})P(BelG/\mathbf{G})P(\mathbf{C}/\mathbf{R})P(\mathbf{G}/\mathbf{R})P(\mathbf{R})}$$

Similarly, from $(Prob_{s'})$ we can derive

$$(Prob_{s'.}) \frac{P(BelC/C)P(BelG/G)P(BelM/M) \sum_{\mathbf{R}} P(C/\mathbf{R})P(G/\mathbf{R})P(M/\mathbf{R})P(\mathbf{R})}{\sum_{\mathbf{C},\mathbf{G},\mathbf{M},\mathbf{R}} P(BelC/\mathbf{C})P(BelG/\mathbf{G})P(BelM/\mathbf{M})P(\mathbf{C}/\mathbf{R})P(\mathbf{G}/\mathbf{R})P(\mathbf{M}/\mathbf{R})P(\mathbf{R})}$$

For definite values of the parameters p , q , s , t , and u , we can compute $Prob_{s.4}$ and $Prob_{s'.4}$. Setting all parameters at .90 yields $Prob_{s'} = P(C, G, M/BelC, BelG, BelM) = .8910 > .7562 = P(C, G/BelC, BelG) = Prob_s$, which is precisely the desired result.

We may conclude that if, under certain independence assumptions, we gain 'coherent' information from highly but not fully reliable witnesses, then our extended belief system may well be more probable than our original belief system, in which case we have a counter-example not only to the Doxastic Extension Principle but also to the Weak Doxastic Extension Principle.

Appendix B

Proof of the Impossibility Theorem

We will consider a case of full agreement between independent reports that are individually credible, while respecting the *ceteris paribus* condition. We will show that there are no informative coherence measures that are truth conducive *ceteris paribus* in such a scenario which I will refer to as a *basic Lewis scenario*. The name is appropriate considering Lewis's reference to relatively unreliable witnesses telling the same story. A number of additional constraints will be imposed on the probabilities involved. The constraints are borrowed from a model proposed by Luc Bovens and his colleagues (2002). That model was in turn devised as an improvement of the model suggested in Olsson (2002*b*). The most salient feature of this sort of model is that the reliability profile of the witnesses is, in a sense, incompletely known. The witnesses may be completely reliable (*R*) or they may be completely unreliable (*U*), and initially we do not know which possibility holds. An interesting consequence of this sort of model is that, from a certain context-dependent level of prior improbability, the posterior probability will be inversely related to the prior: the lower the prior, the higher the posterior. This feature is exploited in the following.

Definition 1: A *basic Lewis scenario* is a pair $\langle \mathbf{S}, \mathbf{P} \rangle$ where $\mathbf{S} = \{ \langle E_1, H \rangle, \langle E_2, H \rangle \}$ and \mathbf{P} a class of probability distributions defined on the algebra generated by propositions $E_1, E_2, R_1, R_2, U_1, U_2$, and H such that $P \in \mathbf{P}$ if and only if:

- (i) $P(R_i) + P(U_i) = 1$
- (ii) $0 < P(H) < 1$
- (iii) $P(E_1/H, R_1) = 1 = P(E_2/H, R_2)$
- (iv) $P(E_1/\neg H, R_1) = 0 = P(E_2/\neg H, R_2)$
- (v) $P(E_1/H, U_1) = P(H) = P(E_2/H, U_1)$
- (vi) $P(E_1/\neg H, U_1) = P(H) = P(E_2/\neg H, U_2)$
- (vii) $P(R_i/H) = P(R_i) = P(R_i/\neg H)$

- (viii) $P(U_i/H) = P(U_i) = P(U_i/\neg H)$
- (ix) $P(E_1/H) = P(E_1/H, E_2)$
- (x) $P(E_1/\neg H) = P(E_1/\neg H, E_2)$
- (xi) $P(R_1) = P(R_2) > 0$

It can be shown that basic Lewis scenarios satisfy the conditions of individual credibility and independence.

Lemma 1: (Theorem 3 in Bovens et al. 2002) Let $\langle \mathbf{S}, \mathbf{P} \rangle$ be a basic Lewis scenario. Letting $h = P(H)$, $\bar{h} = P(\neg H)$, and $r = P(R_i)$,

$$P(H/E_1, E_2) = h^\star = \frac{(h + r\bar{h})^2}{(h + r^2\bar{h})}$$

Lemma 2: (Bovens et al. 2002: 547) Let $\langle \mathbf{S}, \mathbf{P} \rangle$ be a basic Lewis scenario. For all r , h^\star as a function of h has a unique global minimum for $h \in]0, 1[$ which is reached at

$$h_{\min} = \frac{r}{1+r}$$

By calculating the first derivative one can see that h^\star increases (decreases) strictly monotonically for $h > (<) h_{\min}$.

Observation 1: $0 < h^\star < 1$

Observation 2: $h^\star \rightarrow 1$ as $h \rightarrow 0$

Observation 3: $h_{\min} \rightarrow 0$ as $r \rightarrow 0$

Observation 4: $h_{\min} \rightarrow 1/2$ as $r \rightarrow 1$

Definition 2: Let C be a coherence measure. C is *informative* in a basic Lewis scenario $\langle \mathbf{S}, \mathbf{P} \rangle$ if and only if there are $P, P' \in \mathbf{P}$ such that $C_P(\mathbf{S}) \neq C_{P'}(\mathbf{S})$.

Definition 3: A coherence measure C is *truth conducive ceteris paribus* in a basic Lewis scenario $\langle \mathbf{S}, \mathbf{P} \rangle$ if and only if: if $C_P(\mathbf{S}) > C_{P'}(\mathbf{S})$, then $P(\mathbf{S}) > P'(\mathbf{S})$ for all $P, P' \in \mathbf{P}$ such that $P(R_i) = P'(R_i)$.

The stipulation that $P(R_i) = P'(R_i)$ is part of the *ceteris paribus* condition. The other part, concerning independence, is guaranteed already by the fact that we are dealing with Lewis scenarios that, so to speak, have independence built into them.

I will make frequent use in the following of the fact that a probability distribution in \mathbf{P} is uniquely characterized by the probability it assigns to H and R_i . Furthermore, for every pair $\langle r, h \rangle$ there is a probability distribution $P_{r,h}$ in \mathbf{P} such that $P(R_i) = r$ and $P(H) = h$.

Observation 5: $P_{r, h_{\min}(r)}(H/E_1, E_2) \rightarrow 0$ as $r \rightarrow 0$

Impossibility theorem: There are no informative coherence measures that are truth conducive *ceteris paribus* in a basic Lewis scenario.

Proof: We will seek to establish that if C is truth conducive *ceteris paribus* in a basic Lewis scenario, then C is not informative in such a scenario. We recall that the degree of coherence of an evidential system $\mathbf{S} = \{\langle E_1, H \rangle, \langle E_2, H \rangle\}$ is the coherence of the pair $\langle H, H \rangle$. Moreover, if C is a coherence measure then $C(\langle H, H \rangle)$ is defined in terms of the probability of H and its Boolean combinations, as explained in section 6.1 above. In other words, $C_P(\langle H, H \rangle) = C(h)$ where $h = P(H)$. From what we just said it is clear that in order to show that C is not informative, in the sense of $C_P(\mathbf{S}) = C_{P'}(\mathbf{S})$ for all $P, P' \in \mathbf{P}$, it suffices to prove that $C(h)$ is constant for all $h \in]0, 1[$. We will try to accomplish this in two steps, by first showing that $C(h)$ is constant in $I =]0, \frac{1}{2}[$ and then extending this result to the whole interval $]0, 1[$.

Suppose, then, that C is not constant in I . Hence, there are $h_1, h_2 \in I$ such that $C(h_1) \neq C(h_2)$. We may assume $h_1 < h_2$.

Case 1: $C(h_1) > C(h_2)$. By Observation 3, h_{\min} goes to 0 as r goes to 0. Since $h_1 > 0$, it follows that there is a probability of reliability r such that $h_{\min} < h_1$. Consider distributions P_{r, h_1} and P_{r, h_2} in \mathbf{P} . By Lemma 2, h_{\min} is a unique global minimum and h^* is monotonically decreasing for $h > h_{\min}$. Hence, $P_{r, h_1}(h_1/E_1, E_2) < P_{r, h_2}(h_2/E_1, E_2)$. Hence, C is not truth conducive (see Figure B1).

Case 2: $C(h_1) < C(h_2)$. By Observation 4, h_{\min} goes to $1/2$ as r goes to 1. It follows that there is a probability of reliability r such that $h_2 < h_{\min} < 1/2$. Consider distributions P_{r, h_1} and P_{r, h_2} in \mathbf{P} . By Lemma 2, h_{\min} is a unique global minimum and h^* is monotonically increasing for $h < h_{\min}$. Hence, $P_{r, h_1}(h_1/E_1, E_2) > P_{r, h_2}(h_2/E_1, E_2)$. It follows that C is not truth conducive (see Figure B2).




Figure B1: $C(h_1) > C(h_2)$. By choosing r such that $h_{\min} < h_1$ we can construct a counter example to the truth conduciveness of C in the interval $I =]0, \frac{1}{2}[$.

What has been shown so far is that, if C is truth conducive, C is constant in I .

We will proceed to show that, if C is truth conducive, then C is constant in $I' =]\frac{1}{2}, 1[$ as well. Suppose C is truth conducive but not constant in I' . Since C is truth conducive, $C(h) = c$ for all $h \in I$. Since C is assumed not constant in I' , there is an $h \in I'$ such that $C(h) \neq c$.

Case 1: $C(h) > c$.

By Observation 2, $P_{r,h}(H/E_1, E_2)$ goes to 1 as h goes to 0. Since $P_{r,h}(H/E_1, E_2) < 1$, there is a $h' \in I$ such that $P_{r,h'}(H/E_1, E_2) > P_{r,h}(H/E_1, E_2)$, whereas $C(h') = c < C(h)$. This contradicts the assumption of C 's truth conduciveness (see Figure B3).




Figure B2: $C(h_1) < C(h_2)$. By choosing r such that $h_{\min} \in]h_2, \frac{1}{2}[$ we can construct a counter-example to the truth conduciveness of C for $h \in]0, \frac{1}{2}[$.




Figure B3: $C(h) > c$. There is then a point h' such that $C(h') = c < C(h)$ but $P_{r,h'}(H/E_1, E_2) > P_{r,h}(H/E_1, E_2)$.




Figure B4: $C(h) < c$. By choosing r so that $P_{r,h_{\min}}(H/E_1,E_2) < P_{r,h}(H/E_1,E_2)$ we get a counter example to the truth conduciveness of C .

Case 2: $C(h) < c$. By Observation 5, $P_{r,h_{\min}}(H/E_1,E_2)$ goes to 0 as r goes to 0. By Observation 3, h_{\min} goes to 0 as r goes to 0. It follows by these two observations and the fact that $P(h) > 0$ that there is an r such that $P_{r,h_{\min}}(H/E_1,E_2) < P_{r,h}(H/E_1,E_2)$ with $h_{\min} \in I$. Since $h_{\min} \in I$, $C(h_{\min}) = c > C(h)$. We have shown that there is an h' such that $C(h) < C(h')$ and yet $P_{r,h}(H/E_1,E_2) > P_{r,h'}(H/E_1,E_2)$. Again, we have a clash with the assumption that C is truth conducive (see Figure B4).

We have reached a contradiction and may conclude that, if C is truth conducive, then C is constant not only in I but also in I' so that C is in fact constant in the whole interval $]0,1[$. As we said in the beginning, this is sufficient to establish that, if C is truth conducive *ceteris paribus* for a basic Lewis scenario, then C is not informative in such a scenario.

Appendix C

Proofs of Observations

Observation 2.1:
$$P(H/E_1, \dots, E_w) = \frac{1}{1 + (n-1) \left(\frac{1-i}{i(n-1)} \right)^w}$$

in the generalized Huemer model.

Proof:
$$P(H/E_1, \dots, E_w) = \frac{P(E_1, \dots, E_w/H)P(H)}{P(E_1, \dots, E_w)}$$
 (by Bayes's theorem)

$$= \frac{P(E_1/H) \dots P(E_w/H)P(H)}{P(E_1/H) \dots P(E_w/H)P(H) + P(E_1/\neg H) \dots P(E_w/\neg H)P(\neg H)}$$

(by generalized conditional independence)

$$= \frac{i^w \frac{1}{n}}{i^w \frac{1}{n} + \left(\frac{1-i}{n-1} \right)^w \left(\frac{n-1}{n} \right)} = \frac{1}{1 + (n-1) \left(\frac{1-i}{i(n-1)} \right)^w}$$

Observation 3.1: Given (i)–(viii) in section 3.2.3, $P(H/E_1, E_2) = \frac{P(R) + P(H)^2 P(U)}{P(R) + P(H)P(U)}$.

Proof: By Bayes's theorem:

$$(1) \quad P(H/E_1, E_2) = \frac{P(E_1, E_2/H)P(H)}{P(E_1, E_2)}$$

We will now calculate the right-hand side of (1), noting that

$$(2) \quad P(E_1, E_2/H) = P(E_1, E_2/H, R)P(R/H) + P(E_1, E_2/H, U)P(U/H)$$

From (2) and our background assumptions we deduce

$$(3) \quad P(E_1, E_2/H) = P(R) + P(H)^2 P(U)$$

Turning to the denominator of (1),

$$(4) \quad P(E_1, E_2) = P(E_1, E_2/R)P(R) + P(E_1, E_2/U)P(U)$$

By (4) and our assumptions,

$$(5) P(E_1, E_2) = P(R)P(H) + P(H)^2P(U)$$

Finally, by combining (3) and (5) we get (after some simplification)

$$P(H/E_1, E_2) = \frac{P(R) + P(H)^2P(U)}{P(R) + P(H)P(U)}.$$

Observation 3.2: $P(H/E) = P(R) + P(H)P(U)$

Proof: We first note

$$(1) P(E/H) = P(E/H, R)P(R/H) + P(E/H, U)P(U/H) \\ = P(R) + P(H)P(U) \text{ (by (i), (iii), (vii), and (viii))}$$

$$(2) P(E) = P(E/R)P(R) + P(E/U)P(U) \\ = [P(E/R, H)P(H/R) + P(E/R, \neg(H))P(\neg(H/R))]P(R) \\ + [P(E/U, H)P(H/U) + P(E/U, \neg(H))P(\neg(H/U))]P(U) \\ = [P(H)]P(R) + [P(H)P(H) + P(H)P(\neg(H))]P(U) \\ \hspace{15em} \text{(by (i), (iii), (vii), and (viii))} \\ = P(H)P(R) + P(H)P(U)$$

By Bayes's theorem,

$$P(H/E) = \frac{P(E/H)P(H)}{P(E)} \\ = \frac{[P(R) + P(H)P(U)]P(H)}{P(H)P(R) + P(H)P(U)} \text{ (from (1) and (2))} \\ = \frac{P(R) + P(H)P(U)}{P(R) + P(U)} \\ = P(R) + P(H)P(U) \text{ (since } P(R) + P(U) = 1)$$

Observation 3.3: If $P(U)$ is non-extreme, then $P(H/E) > P(H)$.

Proof: By Observation 3.2, $P(H/E) = P(R) + P(H)P(U)$. By algebra, $P(R) + P(H)P(U) > P(H)$ given that $P(U)$ and $P(H)$ are non-extreme.

Observation 3.4: $P(H/E_1, E_2) > P(H/E_1)$.

Proof: Let $a = P(R)$, $b = P(H)$, and $c = P(U)$. We have assumed, as part of the model, that $c < 1$ and $a + b = 1$. The statement to be proved follows from these two assumptions given Observations 3.1. and 3.2. What we need to prove is

$$\frac{a + b^2c}{a + bc} > a + bc.$$

This is established as follows:

$$\begin{aligned}
& \frac{a + b^2c}{a + bc} > a + bc \\
& \Leftrightarrow a + b^2c > (a + bc)^2 \\
& \Leftrightarrow 1 - c + b^2c > (1 - c)^2 + b^2c^2 + 2(1 - c)bc \text{ (by assumption: } a + b = 1) \\
& \Leftrightarrow 1 - c + b^2c > 1 + c^2 - 2c + b^2c^2 + 2bc - 2bc^2 \\
& \Leftrightarrow b^2 - 1 > c - 2 + b^2c + 2b - 2bc \\
& \Leftrightarrow 1 + b^2 - 2b = c(1 + b^2 - 2b) \\
& \Leftrightarrow 1 > c \text{ (by assumption)}
\end{aligned}$$

Observation 3.5: If $P(H/E_1) = P(H/E_2) = P(H)$, then $P(H/E_1, E_2) = P(H/E_1)$.

Proof: We first show that $P(H/E_1) = P(H/E_2) = P(H)$ only if $P(U) = 1$. By Observation 2.2, $P(H/E_1) = P(H)$ only if $P(R) + P(H)P(U) = P(H)$ which, by algebra, entails $P(U) = 1$. Reasoning as in the proof of Observation 3.3, we can now show that $P(U) = 1$ entails

$$\frac{P(R) + P(H)^2P(U)}{P(R) + P(H)P(U)} = P(R) + P(H)P(U)$$

By Observation 3.1, the left-hand side of that equality equals $P(H/E_1, E_2)$ and by Observation 3.2 the right-hand side equals $P(H/E_1)$.

Observation 4.1: Suppose (1) $P(E_i/H) = P(E_i)$, (2) $P(E_1, E_2/H) = P(E_1/H)P(E_2/H)$ and (3) $P(E_1, E_2/\neg H) = P(E_1/\neg H)P(E_2/\neg H)$: Then $P(H/E_1, E_2) = P(H)$.

Proof: Bayes's theorem yields:

$$\begin{aligned}
P(H/E_1, E_2) &= \frac{P(E_1, E_2/H)P(H)}{P(E_1, E_2/H)P(H) + P(E_1, E_2/\neg H)P(\neg H)} \\
&= \frac{P(E_1/H)P(E_2/H)P(H)}{P(E_1/H)P(E_2/H)P(H) + P(E_1/\neg H)P(E_2/\neg H)P(\neg H)} \\
&\quad \text{(By (2) and (3))} \\
&= \frac{P(E_1)P(E_2)P(H)}{P(E_1)P(E_2)P(H) + P(E_1)P(E_2)P(\neg H)} = P(H) \text{ (by (1))}
\end{aligned}$$

Observation 4.2: (Tomoji Shogenji 2002) Suppose that report E lacks individual credibility, so that $P(H/E) = P(H)$. Then $P(L) = (n - 1)P(R)$

and hence $P(L) > P(R)$, when $n > 2$, and $P(L) = P(R)$, when $n = 2$. Moreover, if $P(L) = P(R)$ and $n > 2$, then $P(H/E) > P(H)$.

Proof: If a witness is a truth-teller, her report will be E if and only if H is actually true. If she is a randomizer, she will report E one out of n times no matter what is actually the case. If she is a liar, she will report E only if it is actually the case that H is false; and if it is not the case that H , she tells E one out of $n - 1$ times. Hence,

$$\begin{aligned} P(E) &= P(E/R)P(R) + P(E/U)P(U) + P(E/L)P(L) \\ &= P(E)P(R) + \frac{1}{n}P(U) + [P(E/H,L)P(H/L) \\ &\quad + P(E/\neg H,L)P(\neg H/L)] \\ &= \frac{1}{n}P(R) + \frac{1}{n}P(U) + \left[0 + \frac{1}{n-1}P(L)P(\neg H)\right] \\ &= \frac{1}{n}[P(R) + P(U) + P(L)] \\ &= \frac{1}{n} \end{aligned}$$

The probability that a given witness reports E given that H is true is

$$\begin{aligned} P(E/H) &= P(E/R,H)P(R/H) + P(E/U,H)P(U/H) \\ &\quad + P(E/L,H)P(L/H) \\ &= P(R) + \frac{1}{n}P(U) \end{aligned}$$

Let us now assume that $P(H/E) = P(H)$ or, equivalently, $P(E/H) = P(E)$. It follows that

$$P(R) + \frac{1}{n}P(U) = \frac{1}{n},$$

whence

$$P(U) = 1 - nP(R)$$

But $P(L) = 1 - P(R) - P(U)$, and so

$$P(L) = 1 - P(R) - (1 - nP(R)) = (n - 1)P(R)$$

It follows from (3) that $P(L) = P(R)$, when $n = 2$, and $P(L) > P(R)$, when $n > 2$. By analogous reasoning, that if $P(L) = P(R)$ and $n > 2$, then $P(H/E) > P(H)$.

Observation 4.3: Suppose truth-telling (R), randomization (U), and lying (L) are mutually exclusive and exhaustive hypotheses about the reliability. Then $P(E_1/H, E_2) \approx P(E_1/H)$.

Informal argument: We have

$$\begin{aligned} P(E_1/H, E_2) &= P(E_1/R, H, E_2)P(R/H, E_2) \\ &\quad + P(E_1/U, H, E_2)P(U/H, E_2) \\ &\quad + P(E_1/L, H, E_2)P(L/H, E_2) \end{aligned}$$

and

$$\begin{aligned} P(E_1/H) &= P(E_1/R, H)P(R/H) + P(E_1/U, H)P(U/H) \\ &\quad + P(E_1/L, H)P(L/H) \end{aligned}$$

The ‘liar terms’ in these equations will equal 0 since $P(E_1/L, H, E_2) = 0$ and $P(E_1/L, H) = 0$. Hence,

$$P(E_1/H, E_2) = P(R/H, E_2) + \frac{1}{n}P(U/H, E_2),$$

and

$$P(E_1/H) = P(R) + \frac{1}{n}P(U)$$

Clearly, if the reporter has delivered a true report that fact should raise the probability of her being reliable and diminish the probability of her being a mere randomizer: $P(R/H, E_2) > P(R)$ and $P(U/H, E_2) < P(U)$. As a consequence, we should expect $P(E_1/H, E_2) \approx P(E_1/H)$. We note that it does not matter whether the lying is coordinated or uncoordinated.

Observation 4.4: Suppose truth-telling (R), randomization (U), and lying (L) are mutually exclusive and exhaustive hypotheses about the reliability. Then $P(E_1/\neg H, E_2) \approx P(E_1/\neg H)$, if the liars are uncoordinated. Moreover, $P(E_1/\neg H, E_2) > P(E_1/\neg H)$, if the liars are coordinated and n is large.

Proof: In general,

$$\begin{aligned} P(E_1/\neg H, E_2) &= P(E_1/R, \neg H, E_2)P(R/\neg H, E_2) \\ &\quad + P(E_1/U, \neg H, E_2)P(U/\neg H, E_2) \\ &\quad + P(E_1/L, \neg H, E_2)P(L/\neg H, E_2) \end{aligned}$$

In the case of coordinated lying, the probability of one lying witness’s testifying to the same effect as another lying witness is 1, that is to say, $P(E_1/L, \neg H, E_2) = 1$. Hence,

$$(1) \quad P(E_1/\neg H, E_2) = \frac{1}{n}P(U/\neg H, E_2) + P(L/\neg H, E_2)$$

For uncoordinated lying, one the other hand, $P(E_1/\neg H, E_2) = \frac{1}{n-1}$, and so

$$(2) P(E_1/\neg H, E_2) = \frac{1}{n}P(U/H, E_2) + \frac{1}{n-1}P(L/\neg H, E_2)$$

Now compare each of these two equations with

$$(3) P(E_1/\neg H) = P(E_1/R, \neg H)P(R/\neg H) + P(E_1/U, \neg H)P(U/\neg H) + P(E_1/L, \neg H) \\ = \frac{1}{n}P(U) + \frac{1}{n-1}P(L)$$

By (2) and (3), $P(E_1/\neg(H, E_2)) \approx P(E_1/\neg H)$, if the liars are uncoordinated, since although $P(U) > P(U/H, E_2)$, this will be counteracted by the fact that $P(L/\neg H, E_2) > P(L)$.

It remains to be shown that $P(E_1/\neg H, E_2) > P(E_1/\neg H)$, if the liars are coordinated and n is large. Clearly, (3) goes to 0 as n goes to ∞ . Let us see what happens to (1) as n goes to ∞ . The left-hand term in (1) obviously goes to 0. But what happens to the right-hand term? An application of Bayes's theorem gives

$$P(L/\neg H, E_2) \\ = \frac{P(\neg H, E_2/L)P(L)}{P(\neg H, E_2/L)P(L) + P(\neg H, E_2/U)P(U) + P(\neg H, E_2/R)P(R)}$$

Since $P(\neg H, E_2/R)P(R) = 0$ and $P(\neg H, E_2/U) < P(\neg H, E_2/L)$,

$$P(L/\neg H, E_2) > \frac{P(\neg H, E_2/L)P(L)}{P(\neg H, E_2/L)P(L) + P(\neg H, E_2/L)P(L)P(U)} \\ = \frac{P(L)}{P(L) + P(U)} > 0$$

Hence, while (3) $P(E_1/\neg H)$ goes to 0 as n approaches ∞ , (1) $P(E_1/\neg H, E_2)$ then approaches a constant greater than 0. We may conclude that (1) is greater than (3) if n is large.

Observation 7.1: Suppose that the following hold:

- (i) E_1 and E_2 are independent reports on A_1 and A_2 .
- (ii) $P(A_1/E_1) = P(A_1)$ and $P(A_2/E_2) = P(A_2)$.
- (iii) $A_1 \wedge A_2$, $A_1 \wedge \neg A_2$, $\neg A_1 \wedge A_2$, and $\neg A_1 \wedge \neg A_2$ all have non-zero probability.

Then $P(A_1, A_2/E_1, E_2) = P(A_1, A_2)$.

Proof: By Bayes's theorem,

$$(1) P(A_1, A_2/E_1, E_2) = \frac{P(E_1, E_2/A_1, A_2)P(A_1, A_2)}{P(E_1, E_2)}$$

By conditional independence, $P(E_1, E_2/A_1, A_2) = P(E_1/A_1)P(E_2/A_2)$. It follows from (ii) and familiar probabilistic facts that $P(E_i/A_i) = P(E_i)$, $i = 1, 2$. Hence,

$$(2) P(E_1, E_2/A_1, A_2) = P(E_1)P(E_2)$$

By (iii) and the theorem of total probability, $P(E_1, E_2) = P(E_1, E_2/A_1, A_2)P(A_1, A_2) + P(E_1, E_2/A_1, \neg A_2)P(A_1, \neg A_2) + P(E_1, E_2/\neg A_1, A_2)P(\neg A_1, A_2) + P(E_1, E_2/\neg A_1, \neg A_2)P(\neg A_1, \neg A_2)$. By conditional independence, the right-hand side of that equation equals $P(E_1/A_1)P(E_2/A_2)P(A_1, A_2) + P(E_1/A_1)P(E_2/\neg A_2)P(A_1, \neg A_2) + P(E_1/\neg A_1)P(E_2/A_2)P(\neg A_1, A_2) + P(E_1/\neg A_1)P(E_2/\neg A_2)P(\neg A_1, \neg A_2)$. As already noticed, it follows from (ii) that $P(E_1/A_1) = P(E_1)$ and $P(E_2/A_2) = P(E_2)$. It also follows from (ii) that $P(E_2/\neg A_2) = P(E_2)$ and $P(E_1/\neg A_1) = P(E_1)$. Combining all this yields

$$(3) P(E_1, E_2) = P(E_1)P(E_2)P(A_1, A_2) + P(E_1)P(E_2)P(A_1, \neg A_2) + P(E_1)P(E_2)P(\neg A_1, A_2) + P(E_1)P(E_2)P(\neg A_1, \neg A_2) = P(E_1)P(E_2).$$

It follows from (1), (2), and (3) that $P(A_1 \wedge A_2|E_1, E_2) = P(A_1 \wedge A_2)$, which ends the proof.

Observation 8.1: $P_1(H/E_1) = \langle h^r \rangle = \int_0^1 \frac{h}{h + \frac{h\bar{h}r_1}{h+\bar{h}r_1}} dr_1 = \frac{1+h}{2}$

Proof: By arithmetic simplification,

$$(1) \int_0^1 \frac{h}{h + \frac{h\bar{h}r_1}{h+\bar{h}r_1}} dr_1 = \int_0^1 (\bar{h}r_1 + h) dr_1$$

In general,

$$(2) \int (ax + b)^n dx = \frac{(ax + b)^{n+1}}{(n + 1)a}$$

From (1) and (2),

$$(3) \int_0^1 (\bar{h}r_1 + h) dr_1 = \left[\frac{(\bar{h}r_1 + h)^2}{2\bar{h}} \right]^1 - \left[\frac{(\bar{h}r_1 + h)^2}{2\bar{h}} \right]^0 = \dots = \frac{1+h}{2}$$

Observation 8.2: The function

$$P_2(H|E_1, E_2) = f(h) = \frac{(h + \frac{1}{2}\bar{h})^2}{h + \frac{1}{4}\bar{h}}$$

takes on its minimum for $h = 1/3$ in the interval $h \in (0, 1)$.

Proof: To find the minimum of this function we calculate its derivative with respect to h , set this derivative equal to 0, and solve for $h \in (0, 1)$. By arithmetic simplification and derivation,

$$\frac{d}{dh} \left(h + \frac{1}{2}\bar{h} \right)^2 = \frac{d}{dh} \frac{(h+1)^2}{(3h+1)} = \frac{3h^2 + 2h - 1}{(3h+1)^2}$$

We set the derivate equal to 0 and solve for $h \in (0, 1)$.

$$\frac{3h^2 + 2h - 1}{(3h+1)^2} = 0 \Rightarrow 3h^2 + 2h - 1 = 0 \Leftrightarrow h = \frac{-2 \pm 4}{6}$$

The only extreme value for $h \in (0, 1)$ is $h = 1/3$ which can be verified to be a minimum.

References

- AKIBA, K. (2000). 'Shogenji's Probabilistic Measure of Coherence is Incoherent'. *Analysis*, 60: 356–9.
- ALSTON, W. P. (1993). *The Reliability of Sense Perception*. New York: Cornell University Press.
- BONJOUR, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.
- (1999). 'The Dialectics of Foundationalism and Coherentism', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*. Malden, Mass.: Blackwell, 117–42.
- BOVENS, L., FITELSON, B., HARTMANN, S., and SNYDER, J. (2002). 'Too Odd (not) to Be True: A Reply to Erik J. Olsson'. *British Journal for the Philosophy of Science*, 53: 539–63.
- and HARTMANN, S. (2003). *Bayesian Epistemology*. Oxford: Clarendon Press.
- and OLSSON, E. J. (2000). 'Coherentism, Reliability and Bayesian Networks'. *Mind*, 109: 685–719.
- (2002). 'Believing More, Risking Less: On Coherence, Truth and Non-trivial Extensions'. *Erkenntnis*, 57: 137–50.
- BRANDT, R. B. (1954). 'A Puzzle in Lewis's Theory of Memory'. *Philosophical Studies*, 5: 88–94.
- COADY, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- COHEN, L. J. (1977). *The Probable and the Provable*. Oxford: Clarendon Press.
- (1982). 'What is Necessary for Testimonial Corroboration?' *British Journal for the Philosophy of Science*, 33: 161–4.
- COHEN, S. (1986). 'Knowledge and Context'. *Journal of Philosophy*, 83: 574–83.
- CRAIG, E. (1990). 'Davidson and the Sceptic: The Thumbnail Version'. *Analysis*, 50: 213–14.
- CROSS, C. B. (1999). 'Coherence and Truth Conducive Justification'. *Analysis*, 59: 186–93.
- DAVIDSON, D. (1986). 'A Coherence Theory of Knowledge and Truth', in E. LePore (ed.), *Truth and Interpretation*. Oxford: Blackwell, 307–19.
- DAWID, A. P. (1979). 'Conditional Independence in Statistical Theory'. *Journal of the Royal Statistical Society*, ser. B 41 (no. 1): 1–31.

- DEROSE, K. (1992). 'Contextualism and Knowledge Attributions'. *Philosophy and Phenomenological Research*, 52: 913–29.
- (1995). 'Solving the Skeptical Problem'. *Philosophical Review*, 104: 1–52.
- DEWEY, J. (1929). *The Quest for Certainty: A Study of the Relation of Knowledge and Action*. New York: Capricorn Books.
- DRETSKE, F. (1970). 'Epistemic Operators'. *Journal of Philosophy*, 67: 1007–23.
- EDMAN, M. (1973). 'Adding Independent Pieces of Evidence', in S. Halldén (ed.), *Modality, Morality and Other Problems of Sense and Nonsense*. Lund: Gleerup, 180–8.
- EWING, A. C. (1934). *Idealism: A Critical Survey*. London: Methuen.
- FRANKLIN, J. (1991). 'Healthy Scepticism'. *Philosophy*, 191: 305–24.
- FRAZIER, R. L. (1995). 'Moral Relevance and *Ceteris Paribus* Principles'. *Ratio*, 7: 113–27.
- FRICKER, E. (1995). 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony'. *Mind*, 104: 393–411.
- FUMERTON, R. (1995). Review of Coady (1992). *Philosophical Review*, 104: 618–21.
- GLASS, D. H. (2002). 'Coherence, Explanation, and Bayesian Networks', in Proceedings of the Irish Conference in AI and Cognitive Science, Lecture notes in AI, vol. 2646. New York: Springer-Verlag, 256–9.
- GOLDMAN, A. I. (2001). 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research*, 63: 85–110.
- GRAHAM, P. J. (2000). 'The Reliability of Testimony'. *Philosophy and Phenomenological Research*, 61: 695–709.
- HANSSON, B. (1983). 'Epistemology and Evidence', in P. Gärdenfors, B. Hansson, and N. Sahlin (eds.), *Evidentiary Value: Philosophical, Judicial and Psychological Aspects of a Theory: Essays Dedicated to Sören Halldén on his Sixtieth Birthday*. Lund: Gleerup, 75–97.
- HANSSON, S. O., and OLSSON, E. J. (1999). 'Providing Foundations for Coherentism'. *Erkenntnis*, 51: 243–65.
- HORWICH, P. (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.
- HUEMER, M. (1997). 'Probability and Coherence Justification'. *Southern Journal of Philosophy*, 35: 463–72.
- HUME, D. (1978). *Enquiries Concerning Human Understanding*, ed. L. A. Selby-Bigge. Oxford: Clarendon Press (first published 1777).
- INSOLE, C. J. (2000). 'Seeing Off the Local Threat to Irreducible Knowledge by Testimony'. *Philosophical Quarterly*, 50: 44–56.

- JAMES, W. (1882). 'The Sentiment of Rationality', reprinted 1956 in *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover, 63–110.
- JAMES, W. (1897). 'The Will to Believe', reprinted 1956 in *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover: 1–31.
- JEFFREY, R. (1987). 'Alias Smith and Jones: The Testimony of the Senses'. *Erkenntnis*, 26: 391–9.
- JUSLIN, P., OLSSON, N., and WINMAN, A. (1996). 'Calibration and Diagnosticity of Confidence in Eyewitness Identification: Comments on What Can Be Inferred From the Low Confidence–Accuracy Correlation'. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22: 1304–16.
- KLEIN, P., and WARFIELD, T. A. (1994). 'What Price Coherence?' *Analysis*, 54: 129–32.
- (1996). 'No Help For the Coherentist'. *Analysis*, 56: 118–21.
- LEHRER, K. (1990). *Theory of Knowledge*. 1st edn. Boulder, Colo.: Westview Press.
- (1997). 'Justification, Coherence and Knowledge'. *Erkenntnis*, 50: 243–57.
- (2000). *Theory of Knowledge*. 2nd edn. Boulder, Colo.: Westview Press.
- LESLIE, J. (1989). *Universes*. London: Routledge.
- LEVI, I. (1980). *The Enterprise of Knowledge*. Cambridge, Mass.: MIT Press.
- (1991). *The Fixation of Belief and its Undoing*. Cambridge: Cambridge University Press.
- (2003). 'Contracting from Epistemic Hell is Routine'. *Synthese*, 135: 141–64.
- LEWIS, C. I. (1946). *An Analysis of Knowledge and Valuation*. LaSalle, Ill.: Open Court.
- LEWIS, D. (1979). 'Scorekeeping in Language Games'. *Journal of Philosophical Logic*, 8: 339–59.
- (1996). 'Elusive Knowledge'. *Australasian Journal of Philosophy*, 74: 549–67.
- LUCE, R. D., and RAIFFA, H. (1957). *Games and Decisions: Introduction and Critical Survey*. New York: Wiley.
- LYONS, J. (1997). 'Testimony, Induction and Folk Psychology'. *Australasian Journal of Philosophy*, 75: 163–78.
- MERRICKS, T. (1995). 'On Behalf of the Coherentist'. *Analysis*, 55: 306–9.
- MOORE, G. E. (1925). 'A Defence of Common Sense', in J. H. Muirhead (ed.), *Contemporary British Philosophy*. London: Allen and Unwin, 193–223.

- MOUNCE, H. O. (1997). *The Two Pragmatisms*. London: Routledge.
- NOZICK, R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.
- OLSSON, E. J. (1999). 'Cohering with'. *Erkenntnis*, 50: 273–91.
- (2001). 'Why Coherence is not Truth-Conducive'. *Analysis*, 61: 236–41.
- (2002a). 'What is the Problem of Coherence and Truth?' *Journal of Philosophy*, 99: 246–72.
- (2002b). 'Corroborating Testimony, Probability and Surprise'. *British Journal for the Philosophy of Science*, 53: 273–88.
- (2002c). 'Corroborating Testimony and Ignorance: A Reply to Bovens, Fitelson, Hartmann and Snyder'. *British Journal for the Philosophy of Science*, 53: 565–72.
- (2003a). 'Avoiding Epistemic Hell: Levi on Pragmatism and Inconsistency'. *Synthese*, 135: 119–40.
- (2003b). 'The Epistemology of Keith Lehrer', in E. J. Olsson (ed.), *The Epistemology of Keith Lehrer*. Dordrecht: Kluwer, 1–20.
- (2004). 'Lassen wir den Skeptiker nicht zu Wort kommen: Pragmatismus und radikaler Zweifel', in A. Fuhrmann and E. J. Olsson (eds.), *Pragmatisch denken*. Frankfurt: Ontos Verlag, 247–77.
- and SHOGENJI, T. (forthcoming). 'Can We Trust our Memories? C. I. Lewis's Coherence Argument'. *Synthese*.
- O'NEILL, L. J. (1982). 'Corroborating Testimonies'. *British Journal for the Philosophy of Science*, 33: 60–3.
- PASCAL, B. (1964). *Pensées*, ed. Léon Brunschvicg. Paris: Garnier.
- PEIRCE, C. S. (1877). 'The Fixation of Belief', reprinted 1955 in *Philosophical Writings of Peirce*, ed. J. Buchler. New York: Dover, 5–22.
- (1878). 'The Probability of Induction', reprinted 1955 in *Philosophical Writings of Peirce*, ed. J. Buchler. New York: Dover, 174–89.
- (1905a). 'What Pragmatism Is'. *The Monist*, 161–81.
- (1905b). 'Issues of Pragmatism'. *The Monist*, 481–99.
- (1905c), unpublished manuscript reprinted 1955 in *Philosophical Writings of Peirce*, ed. J. Buchler. New York: Dover, 298–300.
- PLANTINGA, A. (2000). *Warranted Christian Belief*. Oxford: Oxford University Press.
- PRICE, H. H. (1969). *Belief*. London: Allen & Unwin.
- PRITCHARD, D. H. (2002). 'Recent Work on Radical Skepticism'. *American Philosophical Quarterly*, 39: 215–57.
- PUTNAM, H. (1995). *Pragmatism: An Open Question*. Oxford: Basil Blackwell.

- REICHENBACH, H. (1938). *Experience and Prediction*. Chicago: The University of Chicago Press.
- RESCHER, N. (1973). *The Coherence Theory of Truth*. Oxford: Oxford University Press.
- SCHUM, D. A. (1988). 'Probability and the Processes of Discovery, Proof and Choice', in P. Tillers and E. D. Green (eds.), *Probability and Inference in the Law of Evidence: The Uses and Limits of Bayesianism*. London: Kluwer, 213–70.
- SHOGENJI, T. (1999). 'Is Coherence Truth-Conducive?' *Analysis*, 59: 338–45.
- (2001). 'Reply to Akiba on the Probabilistic Measure of Coherence'. *Analysis*, 61: 147–50.
- (2002). 'Justification by Coherence from Scratch', unpublished manuscript.
- SOSA, E. (1999). 'How to Defeat Opposition to Moore'. *Philosophical Perspectives*, 13: 141–54.
- SPOHN, W. (1980). 'Stochastic Independence, Causal Independence, and Shieldability'. *Journal of Philosophical Logic*, 9: 73–99.
- SWAIN, M. (1989). 'BonJour's Coherence Theory of Justification', in J. B. Bender (ed.), *The Current State of the Coherence Theory*. Dordrecht: Kluwer, 115–24.
- TALIAFERRO, C. (1992). 'Imaginary Evil: A Sceptic's Wager'. *Philosophia*, 21: 221–33.
- THAGARD, P. (2000). *Coherence in Thought and Action*. Cambridge, Mass.: MIT Press.
- VAN INWAGEN, P. (1993). *Metaphysics*. Oxford: Oxford University Press.
- WAGNER, C. G. (1991). 'Corroboration and Conditional Positive Relevance'. *Philosophical Studies*, 61: 295–300.
- WEINTRAUB, R. (1997). *The Sceptical Challenge*. London: Routledge.
- WILLIAMS, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Basic Blackwell.

Index

- Akiba, K. 95n1, 101–2
Alston, W. P. 184n11
anomaly 15, 106, 191–2,
Bayesian networks viii, 28, 45–6, 209–10
belief
 and liveliness 183
 cognitively spontaneous 63–76, 92,
 102–5, 136, 206
 coherence justification of 61–76
 common sense 176
 religious 177
BonJour, L. 2–3, 12–13, 17, 23–4, 33,
 90–2, 120, 157n2, 162–5, 173, 184
 anomaly and probabilistic inconsistency
 14–16, 106, 190
 belief, justification of 61–76
 cognitively spontaneous belief 63–76,
 92, 102–5, 136, 206
 coherence and truth 64, 95, 117
 coherence and explanation 110–11
 coherence, concept of 14–16, 20–1,
 95–6, 190
 Doxastic Presumption 20, 61–2, 103
 Observation Requirement 64, 65n3,
 104, 166
 witness agreement 25, 66–9, 134–40
Bovens, L. ix–x, 4, 14n7, 23, 42, 122n4,
 126n6, 140, 211–12
 independence 45–6, 126, 209–10
 impossibility theorem 152–5
 interpretation of Cohen 130–3
 Principle of Indifference 143–50, 152
Brandt, R. B. 152
ceteris paribus clause 3–4, 31–2, 112,
 116–19
Coady, C. A. J. 2, 3, 9n1, 9n2, 11, 24,
 33, 50, 52, 67n5, 91, 146n3, 173
 coherence argument 84–8
 cohesion 84–8
 Hume, argument against 78–81
 individual credibility 81–4, 159
 memory 50–1
 testimony, justification of 77–90
 wager argument 184–6
Cohen, L. J. 2, 4, 38, 40
 corroboration, strength of 128–9
 corroboration theorem 29–30, 47–8,
 60, 75
 independence 27n14, 45, 47, 73
 prior-posterior probability 125–34,
 144–5
Cohen, S. 200
coherence
 alternative measures of 98–100
 concept of 12–16, 14–16, 20–1,
 95–6, 97–8, 100n6, 163–4, 190
 explanatory 108–11, 162–70
 measure, definition of 100–1
 philosophical relevance of 9–12
coherentism 65, 168
cohesion, *see* Coady, C. A.
conditional independence, *see*
 independence
congruence, *see* Lewis, C. I.
consistency
 logical 1, 14, 24, 65n3, 138–9, 157,
 162, 167
 probabilistic 14–16, 190
convergence 26
 conditions for 26–31; *see also* credibility,
 individual; independence
 parameters 31–2, 55–8
corroboration, *see* convergence; *see also*
 Cohen, L. J.
corroboration theorem, *see* Cohen, L. J.
Craig, E. 158
credibility 2–5, 26, 30–1, 157, 166
 and the Principle of Indifference 147–8
 definition of 30
 of cognitively spontaneous beliefs,
 see BonJour, L.
 of memory, *see* Lewis, C. I.
 of testimony, *see* Coady, C. A. J.
Cross, C. B. 95n2, 96n3
 truth conduciveness 104–5
Davidson, D. 4, 12, 52, 67n5, 79, 83, 156
 coherence and truth 157–9
 omniscient interpreter 158
 Principle of Charity 158

- Dawid, A. P. 209
 decision-making
 under risk 144
 under uncertainty 144
 deception 52–5, 69–74, 168–9, 187–95
 DeRose, K. 200, 202
 Dewey, J. 194
 doubt
 external origin of 189
 genuine 49–52, 187–203
 in the credibility of recollections 52–3
 in the credibility of reporters 9–10
 Doxastic Extension Principle 107, 205–10
 Doxastic Presumption, *see* Bonjour, L.
 doxastic system 20, 103, 107–8, 205–10
 Dretske 173, 199
 Dunit example, *see* Klein, P;
 Warfield, T. A.
- Edman, M. 42n4
 evidentiary value 41n4
 Ewing, A. C. vii, 16
 coherence, definition of 13
 coherence as impossible to capture in a
 formula 137–9
- Fitelson, B. x, 42n4, 45n5, 143
 foundationalism 62, 65, 72, 91
 Franklin, J. 175n2
 Frazier, R. L. 117–18
 Fricker, E. 79n4
 Fumerton, R. 78n2
- Glass, D. H. 99
 Goldman, A. I. 29
 Graham, P. J. 79n3
- Hansson, B. 42n4
 Hansson, S. O. 102
 Hartmann, S. x, 42n4, 45n5, 143
 impossibility theorem 152–5
 Principle of Indifference 143–50, 152
 Horwich, P. 36, 190
 surprise, analysis of 38–9
 Huemer, M. 36–7, 48, 59–60, 139, 216
 argument against Bonjour 67–9
 witness model 25–33
 Hume, D. 51, 182n8
 justification of testimony 77–90,
- ignorance 11, 71, 81, 90, 138, 143; *see also*
 Principle of Indifference
- inconsistency, *see* consistency
 incoherence 5, 86, 162–4, 175, 190–3
 independence 2–3, 25, 27–30, 41–2,
 45–8; *see also* Lewis, C. I.
 informational route 85, 187–95
 definition of 85
 impossibility theorem 135; *see also*
 Bovens, L.; Ewing, A. C.;
 Hartmann, S.
 proof of 211–15
 Insole, C. J. 79n3
- James, W. 5, 88n7, 174, 194, 203
 wager argument 176–87
 Jeffrey, R. 27n14
 Juslin, P. 146n2
 justification
 on neutral ground 53, 85, 169, 186,
 192, 193–8
 radical, definition of 12
- Klein, P. x, 2, 3, 95, 104, 112, 116,
 124, 205, 207
 truth conduciveness, argument
 against 105–8
 knowledge 177–8, 198–203
 as the goal of enquiry 194
 as not reducible to reliably acquired
 information 139–40
 common sense 176, 180–1
 limitations of 56
 of non-deception 189–92
 scepticism regarding, *see* scepticism
- Lehrer, K. ix, x, 4, 156, 159–62
 coherence, concept of 160–2
 Leslie, J. 39n2
 Levi, I. ix, 102, 188, 196n17, 201n21
 doubt, reasons for 191–3
 testimony, routine acceptance of
 9n1, 205–6
 Lewis, C. I. x, 2–5, 12, 23, 24, 25, 33,
 61, 62, 65–72, 74–5, 86, 90, 97,
 120, 126–9, 134–6, 139, 143,
 151, 153–4, 162, 164–5, 168, 173
 agreement as leading to practical
 certainty 10–11, 36–7
 basic Lewis scenario, definition of
 211–12
 coherence (congruence), concept of
 13–14, 97–8, 100n6
 independence, definition of 58–60

- memory, credibility of 39–58, 78, 84, 147–9, 159
 memory, justification of 34–58
 Principle of Indifference 151–2
 report, use of the word 11
 supposed facts asserted 17–20, 102, 110, 156–7
 weak foundationalism 65, 72, 91
 Lewis, D. 173, 200–1
 Linda problem 155
 Luce, R. D. 144n1
 Lyons, J. 79n3
- memory
 as an informational route 85
 justification of 3, 34–5, 48–58, 148–52
 report of 10
 Merricks, T. 95n2, 105n8
 Moore, G. E. 174n1, 202
 Mounce, H. O. 177–8, 196n16, 201n21
 wager argument 180–7
- neural network 1, 162–3
 Nozick, R. 199
- Observation Requirement, *see* Bonjour, L.
 Olsson, E. J. 9n1, 14n7, 23, 95n2, 103, 122n4, 130, 133, 140, 161, 162n5, 209, 210, 211
 coherence, as applicable to non-closed sets 102
 coherence measure 99
 witness model 40–2, 144
 Olsson, N. 146n2
 omniscient interpreter 158
 O'Neill 47n6
- Pascal, B. 5, 177, 181n7
 perception 186, 89–90
 as an informational route 85
 Peirce, C. S. 5, 54, 151
 doubt, reasons for 174–5, 187–9, 191
 pragmatism 193–5
 scepticism, reply to 187–204
 Plantinga, A. 184n11
 Price, H. H. 184n12
 Principle of Charity 158
 Principle of Indifference 72n6, 143–50, 152
 Pritchard, D. H. 174n1, 199, 201n20, 202
- probability
 interpretation of 103
 of agreeing testimonies 26, 42, 144
 of a testimonial system 103
 of reliability 37, 43, 127–40
 vs. neural network model 164
 Putnam, H. 187n14
- Raiffa, H. 144n1
 randomization, *see* unreliability
 Reichenbach, H. 184n11
 report, *see* testimony
 reliabilism 139–40
 reliability, *see also* credibility
 as independent of coherence 117–19
 doubtful 10
 independent 45–6
 uncertain (incompletely known) 37, 40–4, 143
 unknown 144–52
 as unjustifiable on neutral ground 195–6
 Rescher, N. 4, 13, 162
 coherence and truth 156–7
 Rescher's Principle 17–18, 206
- scepticism 12, 52, 65, 74–6
 concerning the reliability of witnesses 146
 pragmatist response to 173–204
 Schum, D. A. 42n4
 Shogenji, T. ix, x, 218, 17n9, 70, 95n2
 coherence measure 99–101, 134
 truth conduciveness and strength (specificity) 119–23
 Snyder, J. x, 42n4, 45n5, 143
 Sosa, E. 202n22
 Spohn, W. viii–ix, 209
 supposed facts asserted 17–20, 102, 110, 156–7
 surprise
 analysis of, *see* Horwicz, P.
 as a reason for doubt 189, 191–2
 Swain, M. 13
- Taliaferro, C. 184n10
 testimony
 independent, *see* independence
 in the wide sense 11, 108
 natural 11
 justification of 77–90
 reliable, *see* reliability
 Thagard, P. ix, 4, 24, 111, 156
 coherence and truth 162–70
 coherence, concept of 163–4

- truth conduciveness 102–5, *see also*
 Cross, C. B.; Klein, P.; Shogenji, T;
 Warfield, T.
- unreliability 37, 41–2
- Van Inwagen, P. 39n2
- Wagner, C. G. ix, 48n6
- Warfield, T. A. x, 2, 3, 95, 104,
 112, 116, 124, 205, 207
 truth conduciveness, argument
 against 105–8
- Weintraub, R. 180n6, 184n10
- Williams, M. 200, 201n20
- Winman, A. 146n2