

PROBABILISTIC MEASURES OF COHERENCE

ISBN 90-9019412-6

Cover design: Wendy Bos

Printed by Haveka, Alblasterdam

© Wouter Meijs

All rights reserved. No part of this book may be reproduced in any form or by any means without permission of the author.

Probabilistic Measures of Coherence

Probabilistische maten van coherentie

Proefschrift

**ter verkrijging van de graad van doctor aan de
Erasmus Universiteit Rotterdam
op gezag van de
rector magnificus**

Prof.dr. S.W.J. Lamberts

**en volgens besluit van het College voor Promoties.
De openbare verdediging zal plaatsvinden op**

**donderdag 26 mei om 13:30 uur
door**

Wouter Meijs

geboren te Nuenen

Promotiecommissie

Promotor:

Prof.dr. I.U. Mäki

Overige leden:

Prof.dr. T.A.F. Kuipers

Dr. J.J. Vromen

Prof.dr. R. van Woudenberg

Copromotor:

Dr. I.E.J. Douven

Contents

1	Introduction	1
1.1	Probabilistic Measures of Coherence	1
1.2	Carnap on Concept Explication	2
1.3	Coherence as an Epistemic Virtue	4
1.4	The Coherence Theory of Epistemic Justification	6
1.5	Coherence Between Theory and Evidence	9
1.6	Outline	10
2	Coherence as Mutual Support	15
2.1	Introduction	15
2.2	A Qualitative Theory of Coherence	16
2.3	Quantitative Measures of Coherence	21
2.4	Comparing the Measures C_m	26
2.5	The Extant Measures of Coherence	30
2.5.1	Shogenji's Measure S	30
2.5.2	Olsson's Measure \mathcal{O}	31
2.5.3	Fitelson's Measure C_k	33
2.5.4	Bovens and Hartmann's Difference Function f_r	35
2.6	Anticipated Objections	39
2.6.1	Siebel's Objection Against Fitelson's Measure	39
2.6.2	Bovens and Hartmann's Objection Against Fitelson	41
2.7	Conclusion	42
3	Coherence as Relative Overlap	51
3.1	Introduction	51
3.2	Fitelson's Theory of Coherence	52
3.3	Dependence and Equivalence Desiderata	58
3.4	Measuring Coherence as Relative Overlap	61
3.5	The Extant Measures of Coherence	64
3.6	Coherence as an Epistemic Virtue	66

3.7	A Measure of Epistemic Coherence \mathcal{E}	71
3.8	Conclusion	75
4	Coherence as Confidence Boosting	79
4.1	Introduction	79
4.2	Bovens and Hartmann's Impossibility Result	80
4.3	Bovens and Hartmann's Theory of Coherence	84
4.4	A Different Kind of Impossibility Result	86
4.5	Possible Solutions for Bovens and Hartmann	89
4.6	An Alternative Difference Function	94
4.7	Conclusion	101
5	A Measure of Bootstrap Confirmation	113
5.1	Introduction	113
5.2	Glymour's Qualitative Theory	114
5.3	A Class of Measures of Bootstrap Confirmation	122
5.4	Comparing the Different Measures	128
5.4.1	Threshold Values	128
5.4.2	A Difference Between the Bootstrap Measures	129
5.4.3	Choosing Between the Measures	131
5.5	Anticipated Objections	131
5.5.1	Edidin's and Horwich's Arguments	131
5.5.2	The Problem of Independent Hypotheses	133
5.5.3	Van Fraassen's Critique of Bootstrap Confirmation	134
5.6	Coherence Between Theory and Evidence	137
5.7	Conclusion	139
6	Anticipated Objections	151
6.1	Siebel's Impossibility Argument	152
6.2	Moretti and Akiba's Equi-coherence Principle	154
6.3	Conclusion	160
7	Conclusions	161
	References	165
	Samenvatting	171
	Curriculum Vitae	173

Preface

This thesis comprises the results of my PhD. research in the subject of probabilistic measures of coherence. In recent years, this subject has rapidly gained in popularity and it is the aim of this thesis to make a number of valuable additions to this debate both by proposing a formal framework in which the question of measuring coherence can be assessed and by presenting a number of different measures of coherence. Also, I will critically discuss the proposals that have been made in the literature so far.

A substantial part of this thesis is derived from papers that are either in press or being reviewed at the moment. These papers are:

- Meijs, W. (2005a) 'A Corrective to Bovens and Hartmann's Measure of Coherence,' *Philosophical Studies*, in press.
- Meijs, W. (2005b) 'Coherence as Generalized Logical Equivalence,' manuscript.
- Meijs, W. and I. Douven (2005) 'Bovens and Hartmann on Coherence,' *Mind*, in press.
- Douven, I. and W. Meijs (2005a) 'Measuring Coherence,' *Synthese*, in press.
- Douven, I. and W. Meijs (2005b) 'Bootstrap Confirmation Made Quantitative,' *Synthese*, in press.

Much of my research, but especially the contents of chapters 2 and 5, is the result of joint work with Igor Douven. I would like to thank him for giving me the opportunity to work with him on the subject of measuring coherence and for the enthusiastic and concerned manner in which he has supervised my PhD. research.

My research has greatly profited from discussions with, among others, Luc Bovens, Igor Douven, Branden Fitelson, Stephan Hartmann, Luca Moretti, Jan-Willem Romeyn, Mark Siebel, Jos Uffink, and Janneke van Lith.

Furthermore, I would like to thank those who have commented on one or more of my papers: Luc Bovens, Branden Fitelson, Stephan Hartmann, Franz Huber, Theo Kuipers, Gert-Jan Lokhorst, Luca Moretti, Fred Muller, Bert Postma, Jan-

Willem Romeyn, Mark Siebel, Jos Uffink, Janneke van Lith, Theo van Willigenburg, and Christopher von Bülow.

Also, I would like to thank Michiel Wielema for his very valuable comments on an earlier version of this manuscript and Wendy Bos for designing the cover of this book.

Once, somebody very special said to me that the most important people are the hardest to thank in words. I could not agree more and will let that remark suffice to express a gratitude that words could never describe.

Chapter 1

Introduction

1.1 Probabilistic Measures of Coherence

The concept of coherence has a long history in the philosophy of science and epistemology. A theory that is coherent is by many believed to be a better theory than one that is not, all else being equal. But despite the intuitive interpretation of coherence as ‘hanging together well’ being virtually uncontested, a formal explication has not been forthcoming. Instead, a common complaint by those who refuse to attach any methodological or epistemological value to coherence is that the notion is hopelessly vague. Apparently, this supposedly inherent vagueness was one of the main reasons for Laurence Bonjour to abandon his earlier coherence theory of justification (see, for instance, Bonjour 1997).

The failure of a formal explication of the concept of coherence to materialize is a strong impediment for any epistemological theory in which coherence is an important ingredient. The general object of this thesis is to take up the challenge naturally suggested by this failure by pursuing a number of different explications of coherence. More specifically, I will attempt to explicate some of our intuitive notions of coherence by constructing a number of different probabilistic measures of coherence. Such a project faces various problems. Firstly, it is not clear at all that our intuitive concept of coherence can be captured fully by any formal measure, let alone by a *probabilistic* one. Nevertheless, it seems that the question of the viability of a probabilistic analysis of coherence can only be answered after we have actually engaged in such a project. Only after finding a measure of coherence will we be in a position to test to what extent (if any) it agrees with our intuitive judgements. A second worry is much more important at this stage. For the first problem already presupposes that there is something like a good explication of the concept of coherence. However, it is not clear what the criteria

are by which we can decide that an explication is successful or satisfactory.

Precisely because in general an intuitive concept is rather vague, we cannot exclude the possibility that no explication will fully capture it. Nor can we *a priori* expect there to be only one explication. Therefore, we need a methodological procedure, which will help us arrive at an explication and decide between the different alternatives if there are more than one. This methodological problem will be the subject of this introductory chapter.

In the 40s and 50s of the last century, Rudolf Carnap faced the same type of problems when he tried to formalize the notion of a measure of degree of confirmation. Although I do not share his view that such measures need to be logical in order to be viable, his general methodological remarks will prove to be extremely valuable for the project of measuring coherence.

In this chapter I will start (section 1.2) with Carnap's views on the general requirements for a formal explication of an inexact, prescientific concept. The conclusion from that discussion will be that the description of the informal concept must be very specific in order for its formalization to be as satisfactory as possible. The next three sections, sections 1.3 through 1.5, will discuss three different senses in which the concept of coherence has been used in modern philosophy.

The variety in meaning should be rather unsurprising for a concept that has been around for so long without ever having been given a formal definition. However, it does pose a problem for anyone who attempts to construct a measure of coherence, for it is altogether unclear whether one and the same measure is the best formalization of each of these different senses of coherence. Instead, it will appear in this thesis that different measures of coherence are best suited to deal with some of the different senses in which the concept of coherence has been used in the history of the philosophy of science and in that of epistemology.

1.2 Carnap on Concept Explication

When trying to formalize our intuitive concept of coherence, it is vital to understand first how such a formalization is to be realized and what it means for such a formalization to be successful. These questions are important, especially because they will substantially influence the process by which we arrive at our formalization of coherence. When trying to clarify the concept of degree of confirmation, Carnap clearly saw the importance of a thorough discussion of the methodological requirements for such a project. His account will appear to be very helpful for the project at hand.

In his *Logical Foundations of Probability*, Carnap sets himself the goal of clarifying (1) the concept of degree of confirmation, (2) the logical nature of induction

and (3) the concept of probability (Carnap 1950: 1). According to him, the vital characteristic of the type of clarification that he envisions is that it makes formally precise one of our informal concepts. This is what he calls an *explication*:

[b]y the procedure of *explication* we mean the transformation of an inexact, prescientific concept, the *explicandum*, into a new exact concept, the *explicatum*. ... The explicandum may belong to everyday language or to a previous stage in the development of scientific language. The explicatum must be given by explicit rules for its use, for example, by a definition which incorporates it into a well-constructed system of scientific either logicomathematical or empirical concepts. (*op. cit.*: 3)

The challenging aspect of finding an explication is the fact that the explicandum is not a formal notion and therefore the relation between explicandum and explicatum cannot be one of complete coincidence (*op. cit.*: 5). But how do we then decide whether a certain explication is a *good* explication? According to Carnap, explications may differ with respect to their being *satisfactory*, or, less stringently, in their being more satisfactory than any of the other proposals. For an explication to be satisfactory it must be exact, fruitful (suggesting further research) and as simple as possible, and it must be 'similar to the explicandum' (*op. cit.*: 7). The last requirement is especially ambiguous. If the explicatum does not relate to the explicandum it is useless, but on the other hand, it does not need to be so closely related that in all cases in which the explicandum can be used the explicatum applies also. Precisely because the latter is formally precise, whereas the former is not, substantial differences in applicability may be appropriate or even necessary to avoid inconsistencies.

It appears, therefore, that these requirements are not very restrictive. Many solutions may satisfy them to the same degree or we may disagree about which of them satisfies them best. Furthermore, Carnap proposes no lexical ordering, which implies that when two solutions satisfy different requirements better no decision appears possible. Apparently, more intuitive considerations are necessary to decide which solution is the most satisfactory and we must allow for the possibility that two mutually exclusive solutions are equally satisfactory. Although this is of course an unavoidable consequence of trying to formalize an inexact notion, Carnap emphasizes that we must describe the explicandum as precisely as possible:

There is a temptation to think that, since the explicandum cannot be given in exact terms anyway, it does not matter much how we formulate the problem. But this would be quite wrong. On the contrary, since even in the best case we cannot reach full exactness, we must, in order to prevent the discussion of the problem from becoming entirely futile, do all we can to make at least practically clear what is meant as the explicandum. (Carnap 1950: 4)

To this effect, one may give a general description of the explicandum, or one may provide examples in which the explicatum should be applicable or others in which it should not be applicable. For example, one may indicate that the concept of salt which is to be explicated is not meant in the sense it has in chemistry, but rather in the narrower sense it has in the household language (Carnap 1950: 4–5).

For an explication of the concept of coherence Carnap's remarks are especially important. In common language, 'coherence' has a large number of very different usages. As Luc Bovens and Stephan Hartmann (2003a: 31) note, it makes sense to talk of the coherence of a law firm or an ant-hill, but it seems arguable that the concept of coherence involved is different in both. As a first delineation, I propose to consider coherence in the sense it has been given in the modern philosophy of science literature. In this sense of coherence, coherence is a property of a set of propositions (or beliefs, hypotheses, axioms, sentences, etc.), which can come in degrees: a set can be more or less (in)coherent. Moreover, the coherence of a set has something to do with these propositions 'hanging together'. Nevertheless, this does not yet supply as fully an explanation as possible. For on closer examination it appears that even in the relatively secluded domains of the philosophy of science and epistemology the concept of coherence has been used in a number of (possibly very different) senses. The next three sections will discuss three different (but not necessarily independent) philosophical senses of coherence. My goal is not to do full justice to the discussions in which the concept has emerged, nor to give a full list of all its possible senses. Instead I intend to single out a few senses of coherence, each of which will be formalized in the remainder of this thesis. Naturally, we should not directly exclude the possibility that one measure of coherence is the most satisfactory for all the possible senses of coherence. Nevertheless, given the fact that coherence has functioned in so many different contexts, this seems rather unlikely.

1.3 Coherence as an Epistemic Virtue

The first sense of coherence that this thesis will discuss is coherence as an epistemic virtue. One of the philosophies in which this sense of coherence is well exemplified is Ernan McMullin's version of realism.

In his 'Epistemic Virtue and Theory Appraisal' (1996b), McMullin makes a distinction between a scientific law and a scientific theory. A *law* is empirical in nature and merely reports a correlation. Although this correlation may be either statistical or invariable, it is only the truth of the correlation that a law is concerned with. In contrast, a *theory* gives an account of *why* the correlation holds: '[t]heories attempt to explain why particular regularities recur; they indicate what the causes may be' (McMullin 1996b: 18). And whereas a law is only

concerned with empirical fit, both empirical fit and explanatory success are required to make a good theory. Explanatory success is a vague notion and McMullin instead prefers to present an informal list of virtues that promote a theory's explanatory success. His list of virtues is quite long, and I will not discuss them in depth. Instead, let it suffice to note that coherence is one of the three 'internal' virtues (McMullin 1996b: §4), logical consistency and simplicity being the other two. Interestingly, McMullin defines coherence negatively as the absence of *ad hoc* features, since according to him it is in general the lack of coherence that plays the most important role in the debate. A theory with many *ad hoc* features is said to be much less likely to be true than a theory without such features, *ceteris paribus*.

McMullin is not alone in his view on coherence as a theoretical virtue. Indeed, many others have discussed coherence as one of the virtues that make for good theories. Paul Churchland, for example, lists it as one of the so-called 'superempirical' virtues (next to explanatory power and simplicity) and he, too, argues that a theory's excellence (as measured by the virtue of empirical adequacy combined with the superempirical virtues) 'is the ultimate measure of truth and ontology at all levels of cognition.' (Churchland 1985: 35)

While for McMullin coherence is the absence of *ad hoc* assumptions, many others have given the virtue of coherence a more positive interpretation, in the sense that it is good for a theory if its propositions hang together nicely. For instance, according to Bonjour (1985: 93)

[i]ntuitively, coherence is a matter of how well a body of beliefs 'hangs together': how well its component beliefs fit together, agree or dovetail with each other, so as to produce an organized, tightly structured system of beliefs, rather than either a helter-skelter collection or a set of conflicting subsystems.

Moreover, he believes (*op. cit.*: 95–96) that coherence cannot be equated with logical consistency: clearly a system can be intuitively very incoherent even if it is logically consistent. Arguably, this is what happens if, for instance, two beliefs make each other extremely unlikely without logically contradicting each other. A similar argument (*op. cit.*: 96–97) shows that coherence must be less strict than logical entailment: two beliefs (or propositions) can be very coherent without one of them entailing the other.

Although Bonjour discusses a different sense of coherence – that of coherence as a confidence boosting property, to be discussed in the next section – I believe that his remarks are valuable also for the explication of coherence as an epistemic virtue. Indeed, they seem to come much closer to our intuitive concept of coherence (even if considered only in the sense of coherence as an epistemic virtue) than the mere absence of *ad hoc* hypotheses, especially if we add to them the remark that coherence can come in degrees. However, this description still

remains extremely vague: the descriptions of coherence as hanging together well, as being a matter of degree and as being somewhere in between consistency and entailment in no way determine a unique notion of coherence.

Especially the hanging together intuition seems ambiguous. Indeed, chapters 2 and 3 of this thesis will propose two plausible ways in which this intuition can be explicated. In chapter 2, a view is examined according to which two propositions hang together if they probabilistically support each other and in chapter 3, hanging together is interpreted as set-theoretic overlap. In both cases there are many different measures that can be based on these considerations. However, I will attempt to show that not all of these are equally satisfactory by testing them against some of our intuitions regarding coherence in specific cases.

1.4 The Coherence Theory of Epistemic Justification

Virtually everybody believes that we are justified in holding at least some of our beliefs. Unfortunately, it is much less clear *why* we are justified in holding these beliefs. The problem is that most justifications for our beliefs are beliefs themselves and thus also in need of justification.

For instance, beliefs which cannot be verified by evidential considerations may be justified by beliefs that can be so verified, but how do we justify these evidential beliefs? Some of them may be justified by other evidential beliefs, but how do we justify those beliefs or how can we make plausible that they do not stand in need of justification? According to the infinite regress skeptic, all beliefs need to be justified, but there is no way in which our most basic evidential beliefs can be justified. In the skeptic's view, there are only two options: either the justificatory process will go on and on, or it will circle back on itself, but in either case we have no real justification of our beliefs.

Several solutions to the infinite regress problem have been proposed. Foundationalists believe that there *are* in fact basic beliefs that do not stand in need of any justification. These beliefs provide the foundations for our other beliefs. Although this clearly halts the infinite regress, it is unclear precisely which of our beliefs are basic in the sense of not requiring justification by other beliefs. Coherentists on the other hand believe that there are no such basic beliefs, but that this does not imply that one cannot be justified in holding any beliefs. According to the most popular coherence theory of epistemic justification, one is justified in holding a belief that p iff p is an element of a coherent system of beliefs (cf. Everitt and Fisher 1995: ch. 7, Hetherington 1996: ch. 25).

One of those who have attempted to counter the infinite regress problem by appealing to coherentist considerations is Laurence Bonjour. In his *The Structure of Empirical Knowledge*, Bonjour argues that coherentism is the best theory with

respect to the structure of empirical justification. In his view evidential beliefs are justified because they cohere with our other beliefs: their coherence makes them more likely to be true. Therefore, coherence in this sense can be said to be *truth conducive*. Bonjour even argues that in some cases we may become very sure of a number of quite improbable propositions, if these propositions hang together to a high enough degree. For instance, consider the example by Lewis of a number of ‘relatively unreliable witnesses who independently tell the same circumstantial story’ (Lewis 1946: 346). According to Lewis,

[f]or any one of these reports, taken singly, the extent to which it confirms what is reported may be slight. ... 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. (*ibid.*)

According to Bonjour, this example can even be strengthened in that it is not necessary to assume that the witnesses are reliable at all:

[Lewis’s] 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 much more likely to be false than true (Bonjour 1985: 148, emphasis in original)

The example by Lewis and Bonjour’s consequent remark that no antecedent credibility is required are important from the perspective of probabilistic measures of coherence, since the model of independent and unreliable witnesses has played an important role in many papers on measures of coherence (for example, Klein and Warfield 1994 and 1996; Shogenji 1999; Olsson 2001, 2002 and 2005; Bovens and Hartmann 2003a and 2003b). The basic question that all of these approaches face is whether, given that we are presented with a number of reports by independent and unreliable witnesses, the coherence of the information that they present us with adds to the probability of that information. Here the term ‘witnesses’ may be interpreted very broadly as anything that may report evidence, i.e., other persons, our senses, measuring equipment, etc.

The model in which independent, partially reliable witnesses report on propositions combined with the conception of coherence as being truth conducive forms the basis of the first sense of coherence that this thesis will attempt to

explicate.¹ It is quite plausible that the independent witness-model is not the only option for formalizing the coherentist approach to justification. But given that it has dominated the discussion on probabilistic measures of coherence so strongly, it seems a good place to start. If one were to find that it cannot be formalized in a satisfactory way, one may always try to find other models. In this sense of coherence, then, coherence is inexorably linked to being truth conducive and to a model in which relatively unreliable, independent witnesses supply us with reports to the effect that a certain proposition is true.

Nevertheless, this conception of coherence is quite different from the sense of coherence as an epistemic virtue. As is well-known, McMullin believes explanatory success to be a sign of a theory's truth (McMullin 1996b: 29–31, see also McMullin 1984, 1996a). Yet for him coherence need not be truth conducive in the sense that it has in coherentist accounts of knowledge. Rather, for him coherence is one of the virtues that make a theory more explanatory and, therefore, one of the virtues that together with a number of others makes a theory more likely to be true. Still we should not underestimate the similarities between the two views. Both stress the relation between coherence and the truth of a set of beliefs or propositions. And most of those who have discussed coherence as a confidence boosting property have similarly argued that coherence should be a matter of hanging together and a matter of degree and should be somewhere in between logical consistency and entailment. And until very recently, no attempt had been made to dissolve the vagueness of the concept in the sense of coherence as a confidence boosting property either.

In his (1985), Bonjour still argued that this vagueness was disturbing but not destructive for a theory of coherence, for he believed that coherence is, and seemingly must be, a basic ingredient of virtually all rival epistemological theories as well. For accounts of knowledge of the past and, more generally, of all knowledge that goes beyond the directly empirical, even strongly foundationalist theories will still require an account of coherence. I will come back to this argument in the next section.

However, as announced above, Bonjour later came to repudiate this view and instead adopted a version of the foundationalist theory of justification. One of his reasons for revising his earlier view is precisely the supposedly inherent vagueness of coherence. In his (1999: 124), Bonjour concluded that 'the precise nature of coherence remains a largely unsolved problem' and that this gives strong reason to despair of coherentism with respect to evidential beliefs. Interestingly,

¹Above I have quoted Bonjour's remark that no antecedent credibility is required. Consequently, one may expect that in my model I would allow for the possibility that the witnesses are completely unreliable. However, in my model the reliability of the witnesses will be fixed and for such a model it is the case that if the witnesses are fully unreliable, we cannot become more confident that the information they report is true. They must be reliable to some degree, even if that degree is extremely small. For more on this see Bovens and Hartmann (2003a: 14, 56).

at the time Bonjour drew this pessimistic conclusion, the first probabilistic measures of coherence were being drafted. Moreover, the first of these proposals (Shogenji 1999 and Olsson 2002) were based on Lewis's and Bonjour's claims that coherence can be truth conducive. Therefore, Bonjour's conclusion seems to have been a bit premature and the question whether or not coherence can be truth conducive in the above sense remains an open question.

In chapter 4, I will take up this challenge by discussing critically Bovens and Hartmann's (2003a) work on coherence as a confidence boosting property.

1.5 Coherence Between Theory and Evidence

Above I noted Bonjour's challenge to the foundationalists that they, too, require a theory of coherence to justify theoretical knowledge. In his (1999) he formulates the challenge as follows:

[T]he concept of coherence ... is also an indispensable ingredient in virtually all foundationalist theories; coherence must seemingly be invoked to account for the relation between the basic or foundational beliefs and other nonfoundational or 'superstructure' beliefs, in virtue of which the latter are justified in relation to the former. (Bonjour 1999: 124)

According to Bonjour (1999: 140n15), '[t]he basic point here is that strictly deductive or even enumerative inductive inference from the foundationalist beliefs does not suffice to justify most of the superstructure beliefs that the foundationalist typically wants to claim to be justified.' Without some recourse to the concept of coherence the foundationalist cannot justify the superstructure beliefs by reference to the foundational beliefs. Here the foundational beliefs are similar to what I above called evidential beliefs and the superstructure beliefs are similar to the theoretical beliefs. I will use the latter terminology.

Although I believe that Bonjour is basically right in arguing that the foundationalist cannot do without some concept of coherence, it is not clear that it is the same concept of coherence as any of the two described above (see also Douven 2005b). For example, the independent witness model discussed above does not seem appropriate, especially because there seems to be nothing wrong with the idea that one piece of evidence hangs together well with a theory consisting of a number of propositions (or beliefs or axioms, etc.). We do not need a separate independent source (in this case an evidential belief) to report on each of the different propositions (theoretical beliefs).

Whereas the explications of the other senses of coherence start with our intuitive notion of the respective sense of coherence and try to explicate that notion, in chapter 5 I will not explicitly start out with the sense of coherence as a relation

between the evidence and the theory. Instead, I will attempt to provide a quantitative measure of Clark Glymour's theory of bootstrap *confirmation*. Although I will argue at the end of that chapter that the concept of bootstrap confirmation goes at least some way toward explicating the coherence between a theory and the evidence, I will leave open the question of what changes are required in order to make it a fully satisfactory explication of that notion.

1.6 Outline

In this chapter I have argued that there are at least four different senses in which the concept of coherence has been used in the recent literature on the subject: coherence as a theoretical virtue in the sense of mutual support, coherence as a theoretical virtue in the sense of set-theoretic overlap, coherence as a confidence boosting property, and coherence as a relation between evidence and theory. In each of these cases, coherence has something to do with propositions (or beliefs, axioms, etc.) hanging together and in each of these cases coherence is linked to either truth or justified belief. However, it has become clear from the discussion that the respective senses of coherence also differ in some important aspects. Coherence as used by coherence theories of justification is first and foremost a truth conducive property, especially in the sense that it has been given by many of those who have worked on the subject of measures of coherence. On the other hand, coherence as a theoretical virtue is not a prerequisite for a theory's truth: if a theory has many of the other virtues and a nice empirical fit, its being relatively incoherent may not be all that important. Also, the two different interpretations of this sense of coherence (coherence as mutual support and coherence as relative overlap) seem to point to two very different characteristics of a set. Finally, the coherence between evidential and theoretical statements seems again to be an altogether different relation than the others, and may, therefore, turn out to require an altogether different type of explication than the three above.

As explained above, Carnap believed that after the explicandum has been described or circumscribed as fully as may be, we should try and find an explication that is fruitful, exact, simple and as similar to the explicandum as possible. I believe it will be one of the main lessons of this thesis that these requirements, even if combined with the descriptions of the different senses of coherence above, do not determine the correct measure of coherence for each of the explicanda. Instead, Carnap's (1950: 4) warning that we cannot decide in an exact way whether a certain proposal is right or wrong still applies.

According to John Kemeny and Paul Oppenheim (1952: 308) 'the commonest procedure of explication is to apply a trial and error method till one arrives at an ingenious guess, and then try to find intuitive reasons to justify the proposed ex-

plicatum.' This, they argue, is precisely what Carnap did when, after arriving at a class of symmetric C-functions for his measure of degree of confirmation, he justified his choice for measure C^* by showing that it had satisfactory consequences (*ibid.*). Nevertheless, they feel that such a procedure is 'clearly very dangerous: the intuition of the most honest and well-trained philosopher is likely at times to become a tool for grinding an ax' (*ibid.*). Instead, they propose to lay down a number of desiderata that in combination determine the form of the appropriate measure. For future reference, let us call this the Kemeny–Oppenheim strategy for the explication of a concept.

This strategy will not always be an improvement on the trial and error method, especially in cases where the desiderata are justified by reference to intuitions. Take for example the sense of coherence as mutual support. A quick look at the literature on measures of confirmation reveals that there are a substantial number of strikingly different measures of support (see, for example, Eells and Fitelson (2000) and (2002)), and it may be extremely difficult to find intuitively compelling general desiderata for measures of coherence that significantly decrease the number of potential candidates. In such cases, it might be better to revert to the method of making a few educated guesses and then trying to find out which of these behaves the most satisfactorily in a number of intuitive test cases. In fact, this will turn out to be the best approach for both the concept of coherence as mutual support and for my measure of bootstrap confirmation.

For coherence as relative overlap and for coherence as truth conduciveness, I will argue that there are enough extremely plausible desiderata to determine a single class of measures of coherence.

In this thesis I will explicate the three different senses of coherence as properties of sets of propositions, while the concept of bootstrap confirmation will be explicated as a property of theories, consisting of hypotheses. Nothing important hinges on this of course, but since Glymour has defined bootstrap confirmation as a property of theories, while almost all of those who have proposed measures of coherence have regarded coherence as a property of a set of propositions, I will simply follow suit. Also, the measures I will present will all be completely probabilistic in that they depend only on the probabilistic features of the propositions (hypotheses) in the set (theory) and not on any of their nonprobabilistic features. This is mainly for the reason that it seems otherwise impossible to construct an *exact* explication of coherence or bootstrap confirmation. We must hope that the considerations that often play a role in our intuitive judgements with respect to the notion at hand can be explicated by a purely probabilistic measure.

As a preliminary point, let me also note that throughout much of this thesis I will, like Lewis and most other writers on the subject, simply speak of the coherence of a set of propositions, as if the coherence of such a set were determined solely by the presence or absence of logical relationships between its elements,

or by their contents. But, of course, if coherence is to be given a probabilistic analysis, and if, as may be reasonably assumed (*pace* Carnap 1950), probabilities are not (wholly) a matter of logic, then a set of propositions cannot be coherent or incoherent intrinsically, but only relative to a given probability function defined on the propositions. Thus, the same set of propositions may be coherent relative to one probability function and incoherent relative to another. Concomitantly, a set of propositions may be more coherent than another relative to one probability function, less coherent relative to a second probability function, and the two sets may be equally coherent relative to a third such function. The same remarks apply to my measure of bootstrap confirmation.

Each of the next three chapters will try to explicate a different sense of the concept of coherence. In each chapter it will appear that one (class of) measure(s) is most suited to the specific sense of coherence at stake. Each of the chapters will thus present a different measure of coherence as best capturing the specific sense of coherence under investigation. To forestall misunderstanding, let me emphasize that I am not committed to the argument that these chapters present the most important, or even the most common, senses of coherence. If there are other senses of the concept it would be interesting to find out whether one of the measures proposed in these chapters would be the best explication of it, or whether another measure is more satisfactory. I will come back extensively to this problem in later chapters.

This thesis is organized as follows. Chapters 2 and 3 will discuss two possible explications of the view of coherence as a theoretical virtue. In chapter 2, I will discuss the view of coherence as mutual support and show that there are a potentially very large number of different classes of measures that can be based on this view. By testing some of these classes against an example in which one set is obviously more coherent than the other, I will make a tentative choice for one of these measures. Furthermore, I will show how each of the measures of coherence that so far have been presented is not satisfactory from the viewpoint of coherence as mutual support.

Chapter 3 will explicate a different sense of hanging together, viz., that of coherence as relative overlap. First I will present an intuitive argument why a single measure of coherence that is both a relative overlap measure and a mutual support measure will necessarily be unsatisfactory. Next, I will present a measure that appears to be the best explication of coherence as relative overlap. This measure, too, will be compared with the existing measures.

Chapter 4 will consider coherence as truth conduciveness. This chapter will be based for the largest part on Bovens and Hartmann's (2003a) and (2003b) model of coherence as a confidence boosting property. After having shown that their model is unsatisfactory in several respects, I will propose a model that is very similar to theirs and present a measure that does not suffer from the same

problematic consequences.

Chapter 5 will discuss Glymour's theory of bootstrap confirmation. I will argue that although his version of his theory has been convincingly refuted, these refutations may be evaded by a probabilistic version of his theory. To this end, I will provide a quantitative theory of bootstrap confirmation and present a class of measures that can be considered to measure the degree of bootstrap confirmation. Finally, it will become clear that the resulting measures can also be interpreted as at least partially satisfactory explications of the coherence between a theory and the evidence.

Chapter 6 will discuss a number of anticipated objections to the different measures proposed in this thesis. The recent wave of measures of coherence has been accompanied by an equally impressive stream of arguments against such projects. In this chapter, I will propose a few tentative answers to these criticisms and will attempt to show that none of them can make true their claim that the project should be abandoned altogether.

Chapter 2

Coherence as Mutual Support

2.1 Introduction

In this chapter I will present one explication of the concept of coherence as an epistemic virtue. Many believe that coherence, along with simplicity, consistency and possibly other virtues, is an indicator of a theory's truth and thus distinguishes it from other theories that are possibly equally well confirmed by the evidence but do not exemplify these virtues to the same degree.

In chapter 1, I noted that in this sense of coherence it need not be the case that coherence is truth conducive *per se*. For example, McMullin believes that explanatory success is an indicator of a theory's truth and whereas coherence is one of the contributors to a theory's explanatory success, it is surely not the only one (in his (1996b) he distinguishes at least 10 different virtues) nor is it clear that coherence works on its own and not exclusively in combination with one or more of the other virtues. The problem this poses for an explication of this sense of coherence is that we are only left with the very general remarks that coherence is a matter of degree and that it is a matter of hanging together. And while coherence being a matter of degree can be straightforwardly formalized within a probabilistic framework, the notion of hanging together may be formalized both as mutual support and as set-theoretical overlap. This chapter will explicate coherence as mutual support while the next chapter will consider coherence as relative overlap.

My approach will be to begin as generally as possible by focussing on the 'mutual support' intuition and considering various possible ways of making it precise. This will result in a probabilistic but still qualitative theory of coherence as mutual support (section 2.2). That theory will serve as a stepping stone for defining a family of measures of coherence which seem to capture the notion of

coherence as mutual support equally well (section 2.3). One of the main problems that this chapter encounters is that restricting ourselves to the concept of coherence as mutual support will not uniquely determine which measure of coherence best explicates this concept. For although the discussion in sections 2.2 and 2.3 will exclude a very large number of potential measures, an equally impressive number of candidates remains.

From this it follows that if we are after a single measure of coherence, we will need other considerations to substantially decrease the latter class of measures. This question will be confronted in section 2.4, where I compare each of these measures with an example about which we have very clear intuitions. It will appear that only one of the measures of coherence proposed in section 2.3 satisfies our intuitive judgements in that example. Thus, in parallel to Carnap's case for his measure C^* (see section 1.6), a tentative case can be made for that one measure of coherence on the basis that it has more satisfactory consequences than the others. Next, in section 2.5, the remaining measure is compared with the extant accounts of coherence and it is argued that it is preferable as a measure of coherence as mutual support.

Finally, I consider and seek to defuse two objections to which the preferred measure (as well as the other measures proposed in section 2.3) seems vulnerable (section 2.6).

2.2 A Qualitative Theory of Coherence as Mutual Support

Almost all the extant accounts of probabilistic measures of coherence place great emphasis on the notion of coherence as mutual probabilistic support. Tomoji Shogenji (1999) and Branden Fitelson (2003) have both discussed coherence explicitly as a measure of mutual support, while Bovens and Hartmann (2003a) give mutual support a central place in their theory of coherence, too. Presumably, the first to make the link between coherence and mutual support was C. I. Lewis, who defined coherence (or congruence, as he called it) for a set of statements as follows:

A set of statements ... will be said to be congruent [i.e., coherent] 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.
(Lewis 1946: 338)

Lewis's definition clearly explicates coherence as probabilistic mutual support. However, it also faces a number of difficulties. Firstly, it is still a qualitative theory. As noted in chapter 1, intuitively coherence can come in degrees, which

means that a satisfactory explication of coherence should allow for more than mere categorical (yes/no) verdicts about the coherence of sets of propositions. Secondly, Lewis's theory only considers the support a single proposition receives from all the others, while there are clearly many other relations between the propositions in a set that may matter for its coherence. In this section, I will show how his qualitative framework can be modified in order to set to rights this shortcoming, while the next section shows how his revised theory can be adapted so as to make coherence a matter of degree.

The reason why Lewis's theory does not seem to capture adequately our notion of mutual support is that it only considers the (in)dependence of each single *one* element of a set on *all* the remaining ones. This I will call *one-all coherence*:

Definition 2.1 A set $S = \{R_1, \dots, R_n\}$ is one-all coherent *precisely if for all* $i \in \{1, \dots, n\}$ *it holds that* $p(R_i | R_1 \wedge \dots \wedge R_{i-1} \wedge R_{i+1} \wedge \dots \wedge R_n) > p(R_i)$.

However, intuitively, other relations between propositions may equally add to a set's propositions being mutually supportive. Therefore, this definition should be strengthened. A first obvious strengthening is to consider the dependence of each *one* element on *any* non-empty subset not including that element.¹ Formally:

Definition 2.2 A set $S = \{R_1, \dots, R_n\}$ is one-any coherent *precisely if for all* $i \in \{1, \dots, n\}$ *and all* $S' \subset S \setminus \{R_i\}$ *such that* $S' \neq \emptyset$ *it holds that* $p(R_i | \bigwedge S') > p(R_i)$.

A similarly obvious strengthening is to consider, for each *partition* of a set into two non-empty subsets, the dependence of these subsets on one another, resulting in²

Definition 2.3 A set $S = \{R_1, \dots, R_n\}$ is partition coherent *precisely if for all non-empty* $S', S^* \subset S$ *such that* $S' \cap S^* = \emptyset$ *and* $S' \cup S^* = S$ *it holds that* $p(\bigwedge S' | \bigwedge S^*) > p(\bigwedge S')$.

And the foregoing definitions can be further strengthened by considering the dependence of *any* non-empty subset on *any* other non-overlapping, non-empty subset, as follows:

¹Another possible strengthening is to require not only that each element be supported by the conjunction of the other elements but also that each element have a probability above a given threshold value conditional on the conjunction of the other elements (this would go some way toward formalizing Hage's (2004) conception of coherence, which requires, among other things, that each element of a set be rationally acceptable relative to the conjunction of the other elements of the set). Definitions 2.2, 2.3, and 2.4 could be strengthened in the same way.

²The notion of partition coherence was suggested to Igor Douven and me by an anonymous referee for *Synthese*, for which we are very grateful.

Definition 2.4 A set $S = \{R_1, \dots, R_n\}$ is any-any coherent *precisely if for all non-empty* $S', S^* \subset S$ *such that* $S' \cap S^* = \emptyset$ *it holds that* $p(\bigwedge S' \mid \bigwedge S^*) > p(\bigwedge S')$.

Before continuing, let me verify that the definitions above are genuinely distinct notions. It is evident that any-any coherence implies both one-any coherence and partition coherence, and that both one-any coherence and partition coherence imply one-all coherence. However, the converse implications do *not* hold. More specifically (for a proof, see Appendix 2 A)³

Theorem 2.1

1. *One-all coherence does not imply one-any coherence.*
2. *One-all coherence does not imply partition coherence.*
3. *One-any coherence does not imply any-any coherence.*
4. *One-any coherence does not imply partition coherence.*
5. *Partition coherence does not imply any-any coherence.*
6. *Partition coherence does not imply one-any coherence.*

It follows from theorems 2.1.4 and 2.1.6 that one-any coherence and partition coherence are logically independent notions. This means that there is even a fifth category of coherence:

Definition 2.5 A set S is one-any + partition coherent *precisely if it is both one-any coherent and partition coherent.*

While any-any coherence obviously implies one-any + partition coherence, I have so far not succeeded in proving that the converse does not hold. Nevertheless, I feel confident enough to make

Conjecture 2.1

One-any + partition coherence does not imply any-any coherence.

Figures 2.1 and 2.2 give a graphic representation of the results so far, given that the conjecture does or does not hold, respectively.

Although we are after a probabilistic measure of coherence and not of incoherence, it seems helpful to define incoherence also. In the case of incoherence, we will have the distinct notions of one-all, one-any, partition, and any-any *in*-coherence. These can simply be obtained by replacing ‘>’ by ‘<’ in definitions 2.1–2.4 (one proves that these are genuinely distinct notions of incoherence in the same way as was done for the corresponding notions of coherence).

³The proofs of this and various other theorems in this thesis have been derived with the help of the PrSAT package for *Mathematica* (version 5), which was written by Branden Fitelson and Jason Alexander and is publicly available via <http://www.fitelson.org>; for some of the theory behind it, see Fitelson (2001a: 93–100).




Figure 2.1: The relations between the different definitions of coherence, if conjecture 2.1 is true. ‘O-A + Part’ stands for One-Any + Partition Coherence

Assuming that the conjecture is true, these definitions yield a qualitative theory of coherence which divides the class of all finite sets of propositions into eleven categories, namely, those that are only one-all coherent, those that in addition are one-any coherent but not partition coherent, those that are partition coherent but not one-any coherent, those that are both one-any and partition coherent but not any-any coherent, those that are any-any coherent, those that are only one-all incoherent, and so on, plus, of course, a category ‘neither of the foregoing.’ The categorization might be extended by defining various other notions of probabilistic independence for sets of propositions, most obviously by defining different relations of coherence neutrality (for example, by replacing ‘>’ by ‘=’ in definitions 2.1-2.4). Below, I will not need such different notions of neutral coherence, therefore it will suffice to define a set as being *independent* by

Definition 2.6 A set $S = \{R_1, \dots, R_n\}$ is independent precisely if for all non-empty $S', S^* \subset S$ such that $S' \cap S^* = \emptyset$ it holds that $p(\bigwedge S' \mid \bigwedge S^*) = p(\bigwedge S')$.

In order to enhance readability, I will no longer use the ‘ \bigwedge ’-symbols when discussing probabilities in the rest of this thesis. Instead, it will simply be assumed that if H is a proposition, then $p(H)$ constitutes the probability of that proposition (relative to the background knowledge, reference to which will likewise be suppressed in the rest of this thesis) and if S is a (finite) set of propositions then




Figure 2.2: The relations between the different definitions of coherence, if conjecture 2.1 is false

$p(S)$ constitutes the probability that the conjunction of the propositions in that set (i.e., $\bigwedge S$) is true (again, relative to the background knowledge). And similarly for $p(\neg S)$, $p(H \wedge S)$, which constitutes the probability that $H \wedge (\bigwedge S)$ is true, $p(S|S')$, etc.

It seems that the notion of any-any coherence (definition 2.4) provides the best qualitative explication of the notion of coherence as mutual support. After all, one-all coherence may hold for a set $S = \{R_1, \dots, R_n\}$ and yet, as can be gleaned from the proof of theorem 2.1.1, there may be $i, j \in \{1, \dots, n\}$ such that $p(R_i | R_j) < p(R_i)$, that is, it may contain propositions that undermine one another and thus do the exact opposite of being mutually supportive. And although there can be no single propositions undermining each other in a set that is one-any coherent, it is an entirely straightforward, albeit somewhat tedious, task to verify that it follows from the second probability model specified in the proof of theorem 2.1 that a set $S = \{R_1, \dots, R_n\}$ can be one-any coherent and yet for all $i, j, k, l \in \{1, \dots, n\}$ such that $i \neq j \neq k \neq l$ it may hold that $p(R_i \wedge R_j | R_k \wedge R_l) < p(R_i \wedge R_j)$.⁴ Arguably, if each conjunction of two propositions in a set is undermined by any conjunction of two other propositions in that set, then the propositions in the set do not really hang together all that well.

⁴A point concerning notation: throughout this thesis, I use $i \neq j \neq k$ as an abbreviation of $i \neq j \wedge j \neq k \wedge i \neq k$, and similarly for similar expressions.

A similar claim holds for the notion of partition coherence: the third probability model specified in the proof of theorem 2.1 shows that a set of propositions may be partition coherent, yet for each proposition in the set there may be a pair of other propositions in the set whose conjunction undermines the former. If, on the other hand, a set is any-any coherent, then every ‘bit’ of it is supported by every other ‘bit’ of it. It seems that the propositions of such a set can be truly said to be mutually supportive.

This concludes the qualitative explication of coherence as mutual support. As noted above, to be satisfactory an explication of the concept of coherence should be sensitive to different *degrees* of coherence. This is especially important, since the sensitivity to degree may even interfere with the different qualitative classes distinguished above. For instance, it may not be the case that each any-any coherent set of propositions will intuitively always be more coherent than a one-any or even a one-all coherent set. In a one-any coherent set, at least some parts of the set may undermine each other with respect to some other parts, but this undermining may be compensated for by other parts’ supporting each other very well. In that case, such a set may turn out to be intuitively more coherent than an any-any coherent set, in which every ‘bit’ of it is only minimally supported by any of its other ‘bits’. Therefore, we need our explication to be sensitive to the degrees of mutual support. I now turn to the task of generalizing the above qualitative theory of coherence into a quantitative theory of coherence as mutual support.

2.3 Quantitative Measures of Coherence as Mutual Support

This section will begin with showing how the above general qualitative framework can be generalized toward a general quantitative framework of coherence as mutual support. The result will be five different classes of measures of coherence, each of which in turn consists of a large number of different classes of measures. With the help of two plausible considerations it will appear possible to limit significantly the number of potential measures of coherence as mutual support. However, the conclusion of this section will still be that the sense of coherence as mutual support allows for an important number of strikingly different measures.

Each of the above qualitative classes of coherence can readily be made quantitative by using one of the Bayesian measures of support (which will be introduced below). Recall that one-all coherence is a matter of each proposition in a set being probabilistically supported by the conjunction of all the other propositions in that set. To obtain a corresponding measure of coherence one can simply

proceed as follows: measure the *extent* to which each proposition is supported by the conjunction of the remaining members of the set and take some average of the results for all propositions in the set. The same procedure can be followed for the other four notions of coherence.

To make this idea more precise, let m be some measure of confirmation and let $[S]$ indicate the set of ordered pairs of non-empty non-overlapping subsets of $S = \{R_1, \dots, R_n\}$, that is, $[S] = \{\langle S', S^* \rangle \mid S', S^* \subset S \setminus \emptyset \wedge S' \cap S^* = \emptyset\}$. Then, in correspondence to definitions 2.1–2.5, let us define the following sets:

$$\alpha(S) = \{m(R_i, S \setminus \{R_i\}) \mid 1 \leq i \leq n\} \quad (2.1)$$

$$\beta(S) = \{m(R_i, S') \mid 1 \leq i \leq n \wedge S' \subset S \setminus \{R_i\}\} \quad (2.2)$$

$$\gamma(S) = \{m(S', S^*) \mid \langle S', S^* \rangle \in [S] \wedge S' \cup S^* = S\} \quad (2.3)$$

$$\delta(S) = \beta(S) \cup \gamma(S) \quad (2.4)$$

$$\epsilon(S) = \{m(S', S^*) \mid \langle S', S^* \rangle \in [S]\} \quad (2.5)$$

These definitions provide us with five different classes of measures of coherence, to wit, those that take the degree of coherence of a given set S to be some average of $\alpha(S)$, $\beta(S)$, $\gamma(S)$, $\delta(S)$ and $\epsilon(S)$, respectively. Each of these classes contains infinitely many subclasses of measures, since there are in principle infinitely many weighing procedures for each of these classes. Moreover, each of these subclasses again has a number of subclasses, since so far more than ten measures of confirmation have been proposed in the literature. Thus, it will be clear that each of these classes contains a vast number of different measures of coherence.⁵

How can we commence decreasing this vast number of different measures of coherence? Well, for a start it certainly seems preferable to have a measure of coherence that takes into account *all* the correlations and anti-correlations that may exist between the different non-empty subsets of a set of propositions, i.e., it seems preferable to measure the coherence of a set S of propositions by taking some average of $\epsilon(S)$. A measure based on $\beta(S)$ would for instance miss the negative dependencies between all pairs of conjunctions of two propositions not sharing a proposition in the case of the set of propositions specified in the proof of theorem 2.1.3; a similar remark would apply to a measure based on $\gamma(S)$.⁶ Furthermore, absent an argument to the contrary, one may argue that a straight

⁵And yet these classes of measures are not exhaustive; for instance, both Shogenji's (1999) measure and Olsson's (2002) measure – see the next section for definitions – are outside any of these classes.

⁶Fitelson (2003) rightly criticizes Shogenji's (1999) measure of coherence for only considering whether the set as a whole is dependent or not. However, Fitelson's own measure as presented in his (2003) is open to the same type of criticism, given that it does not take into consideration correlations that may exist between some of the subsets of a set. More recently, Fitelson (2004) has reformulated his measure in order to make it sensitive to relations of the latter kind, too.

average seems preferable to a weighted one.⁷ Although these considerations narrow considerably the number of remaining candidates – only the ten or so different measures of support remain – it is much less clear on what grounds we are to choose one of these measures as our measure of coherence as mutual support.

In section 1.6, I distinguished two different approaches for arriving at an explication of an intuitive concept. One of these – presumably the one Carnap availed himself of – involves making a few educated guesses on the basis of the explanations and examples presented to describe the intuitive concept. Out of these educated guesses, the one with the most favorable consequences is put forward as the most satisfactory explication. The other approach, proposed by Kemeny and Oppenheim's (1952), is an attempt to construct as many desiderata as required to narrow down the class of viable explications enough to single out one alternative. Above I have indicated that while the latter approach, if feasible, is the most satisfactory, it is not at all clear that we can always find (enough) desiderata that are intuitively compelling. This seems to be the case for a measure of coherence as mutual support. For which requirements could be proposed in order to decide between the various measures of coherence? In his (2003), Fitelson proposes that sets consisting of only equivalent propositions are always more coherent than sets that do not consist of only equivalent propositions. Unfortunately, as will be shown in the next chapter, the combination of this requirement with a view of coherence as mutual support leads to some very counterintuitive results. Since until now no other general desideratum has been proposed for measures of coherence, we thus have no other choice than to pursue a few educated guesses and see whether at least one of these has satisfactory consequences.

The question of whether a measure has satisfactory consequences will be answered by presenting a number of *test cases*, that is, cases in which there is an intuitively indisputable verdict as to which of two sets of propositions is the more coherent one (*if* there is a more coherent one; sets can intuitively be equally coherent of course).⁸ And as it will turn out, at least one of the measures that is proposed below is able to handle perfectly the test cases, those already present in the literature as well as those to be added in the present chapter, and in effect is, on balance, better able to handle these than any of the measures currently to be found in the literature.

Let me begin with narrowing down the number of possible measures to the three which seem to be the most popular.⁹ Evidently, if all of these turn out to be

⁷But given such an argument to the contrary, it is straightforward to adapt the measures of coherence to be defined below to ones that take weighted averages.

⁸Bovens and Hartmann (2003a, Ch. 2) is a real treasure trove for such test cases.

⁹For a discussion and comparison of these measures, see Eells and Fitelson (2002), who make a strong case for *d* and *l* on the basis of symmetry considerations.

unsatisfactory, we can still turn to one of the others and see how these behave.¹⁰ The three that I will consider in this chapter are:

- *the difference measure*: $d(R_1, R_2) =_{df} p(R_1 | R_2) - p(R_1)$;
- *the (log-)ratio measure*: $r(R_1, R_2) =_{df} p(R_1 | R_2) / p(R_1)$ (or, as some prefer, $r(R_1, R_2) = \log [p(R_1 | R_2) / p(R_1)]$);
- *the (log-)likelihood measure*: $l(R_1, R_2) =_{df} p(R_2 | R_1) / p(R_2 | \neg R_1)$ (or a logarithm of that ratio).

It is readily seen how each of these can be generalized so as to apply to ordered pairs of sets. Let us denote the generalized versions of the difference measure, ratio measure and likelihood measure as d , r and l , respectively.

- $d(\langle S, S' \rangle) =_{df} p(S | S') - p(S)$;
- $r(\langle S, S' \rangle) =_{df} p(S | S') / p(S)$;
- $l(\langle S, S' \rangle) =_{df} p(S' | S) / p(S' | \neg S)$.

Let m be a variable for such generalized measures. Further, let $\llbracket S \rrbracket$ denote the number indicating the cardinality of S .¹¹ Then the following defines a family of three measures of coherence as relative overlap:

Definition 2.7 *Given a set $S = \{R_1, \dots, R_n\}$ and an ordering $\langle \hat{S}_1, \dots, \hat{S}_{\llbracket S \rrbracket} \rangle$ of the members of S , the degree of m -coherence of S is given by the function*

$$C_m(S) =_{df} \frac{\sum_{i=1}^{\llbracket S \rrbracket} m(\hat{S}_i)}{\llbracket S \rrbracket}, \quad (2.6)$$

for $m \in \{d, r, l\}$.

The question of choosing between d , r and l would be substantially relieved if all of the resulting measures of coherence were ordinally equivalent. Here, two

¹⁰ Other measures of coherence can be found in, for example, Kemeny and Oppenheim (1952), Nozick (1981), Gaifman (1985), Christensen (1999), Eells and Fitelson (2000), and Kuipers (2000). Note that if a measure of confirmation has been shown to be problematic or even unsuited *qua* measure of confirmation, then that does not automatically mean it is unsuited to build a measure of coherence on. I will come back extensively to the Kemeny–Oppenheim measure in chapter 3.

¹¹ Note that for any set $S = \{R_1, \dots, R_n\}$ it holds that $\llbracket S \rrbracket = \sum_{i=1}^{n-1} \binom{n}{i} (2^{n-i} - 1)$. After all, from every set with n elements we can choose $\binom{n}{1}$ singletons as the first element of an ordered pair, $\binom{n}{2}$ sets of two elements as the first element of an ordered pair, etc. For each subset with i elements ($i \leq n$) there exist $2^{n-i} - 1$ non-empty non-overlapping subsets. So, for every subset with i elements that serves as the first member of an ordered pair we can choose $2^{n-i} - 1$ subsets of the required sort that can serve as a second member.

measures of coherence m and m' are ordinally equivalent if, for all sets S, S' , $m(S) > m(S')$ iff $m'(S) > m'(S')$ (cf. Fitelson 2005b). If all measures of coherence always agreed on which set is the more coherent of two sets of propositions, then our measure of coherence would at least get all the comparative judgements right. Moreover, to arrive at the most satisfactory measure, we could simply compare each of these measures with different sets about which we have clear intuitions with respect to the degree of their coherence. However, as the next theorem shows, things are not that easy (see Appendix 2 B for a proof):

Theorem 2.2 *For any pair of distinct $m, m' \in \{d, r, l\}$ there exist infinitely many triples $\langle S, S', p(\cdot) \rangle$ such that either $C_m(S) > C_m(S')$ and $C_{m'}(S) < C_{m'}(S')$ or $C_m(S) < C_m(S')$ and $C_{m'}(S) > C_{m'}(S')$.*

This proposition does not guarantee at all that the measures also disagree in any of the test cases of coherence. That is, there may exist no triples $\langle S, S', p(\cdot) \rangle$ such that the measures do not all agree on which of the sets is the more coherent one (relative to $p(\cdot)$) and such that one set is *intuitively* clearly more coherent than the other. In fact, when applied to the many test cases that so far have been presented in the literature, all measures C_m assign a higher value to the set that is intuitively the more coherent one. Nonetheless, in the next section I will compare the three different measures defined above in the light of some new, relatively clear examples, and we will see that there is at least some reason to prefer C_d to the other measures.

Preliminary to that, let me briefly note that in the presence of a quantitative theory of coherence, the qualitative theory of section 2.2 is not altogether redundant. For although – as noted above – it may very well be the case that an any-any coherent set is intuitively less coherent than a one-any coherent set, the definitions of any-any coherence and any-any incoherence are quite valuable as limiting cases of coherence as mutual support. That is, if a set is any-any coherent, then it should have positive coherence and if it is any-any incoherent, it should have negative coherence. Thus, we have a good guideline for judging measures with respect to their being measures of coherence as mutual support: if they do not assign all any-any coherent sets both a positive degree of coherence and a higher degree of coherence than all independent sets have, then they cannot be measures of coherence as mutual support.

Moreover, in combination the two theories suggest some interesting psychological questions. To mention but one, it would be interesting to know whether people are more likely to believe a story that is any-any coherent than one that is only one-any coherent even if the second has a higher degree of coherence than the first. This question and similar ones our theories allow us to formulate seem open to empirical investigation.

2.4 Comparing the Measures C_m

Above I indicated that the view of coherence as mutual support does not determine which of the many measures of confirmation is the correct one for the project of measuring this type of coherence. Moreover, I also noted that there does not seem to be a general requirement that could be combined with the view of coherence as mutual support in order to significantly decrease the number of measures of coherence left. Therefore, it seems appropriate to test each of the measures that have been presented so far against some of our more specific intuitions about coherence. If one of the measures performs better in all the test cases of coherence, then it seems that this measure is the most satisfactory. A test case in this context is an example in which two or more sets are compared with each other and about which we intuitively feel very strongly which of the two is more coherent. Or, alternatively, a test case is an example about which a group of respondents would almost unanimously agree which one of the two (or more) sets presented in the example is the most coherent.

In this chapter, coherence is explicated in terms of mutual support. Unfortunately, it seems too far-fetched to argue that we have specific intuitions with respect to this (broad) notion of coherence as mutual support. Instead, we must revert to our basic intuitions with respect to the general notion of coherence. This need not be problematic. As argued in chapter 1, all the different senses of coherence share the notion of ‘hanging together,’ and it may turn out that in many cases this notion is defined well enough for us to make an intuitive judgement that should be respected by all different explications of the concept of coherence. Indeed, in this and the next chapters it will appear that the majority of test-cases are in fact respected by all the preferred explications of the different senses of coherence.

On the other hand, we should not always expect this to be the case. If the different senses of coherence are best captured by different explications then we should expect that some of our intuitions are only compatible with some of these explications. For instance, as we will see below (section 2.6), some of the counterexamples that have been proposed against Fitelson’s measure of coherence as mutual support are based on a conception of coherence as relative overlap.

In such cases, the example cannot be a genuine test-case in the sense that it is a test case for all different explications of coherence. Be that as it may, in this section I will present an example about which – I will argue – we have clear intuitions with respect to coherence and which certainly seems applicable to the project of explicating coherence as a theoretical virtue (and, *a fortiori*, to the sense of coherence as mutual support). More specifically, I will argue that we feel intuitively that one of the two sets I will describe below is much more coherent than the other. It will appear that only C_d supports this intuition. Thus,

this example shows that measure C_d has more satisfactory consequences than the other two and, therefore, that we have a reason to (weakly) prefer it over the others, pending any examples to the contrary.¹²

Here is the example. Consider two situations in both of which it is assumed that a murder has been committed in a street in a big city with 10,000,000 inhabitants, 1,059 of them being Japanese, 1,059 of them owning Samurai swords, and 9 of them both being Japanese and owning Samurai swords. In situation I we assume that the murderer lives in the city and that everyone living in the city is equally likely to be the murderer. In situation II, on the other hand, we make the assumption that the victim was murdered by someone living in the street in which her body was found. In that street live 100 persons, 10 of them being Japanese, 10 owning a Samurai sword, and 9 both being Japanese and owning a Samurai sword.




Figure 2.3: Diagrams of the probability distributions corresponding to situations I (left) and II (right)

Next, consider the set of propositions $S = \{J, O\}$ with

- J : The murderer is Japanese;
- O : The murderer owns a Samurai sword.

Let us compare the degrees of coherence assigned to S by the three different measures of coherence as mutual support, when we make the probabilistic assumptions corresponding to situation I and, respectively, situation II (see figure 2.3;

¹²One may feel that there is also a *strong* reason to prefer C_d over C_l . For – as can easily be seen – C_l is not defined if a set of propositions has at least two subsets S' and S^* such that $p(S' \mid \neg S^*) = 0$. However, in chapter 6, I will argue that there are good reasons to exclude such sets entirely from the intended domain of our measures of coherence.

note that the diagrams are not drawn to scale). These are given in the following table:

	I	II
C_d	.0084	.8
C_r	80.3	9.0
C_l	80.9	81.0

What does intuition say about these situations? Clearly, an appeal to mutual support considerations will not be very helpful here. In both situations, propositions J and O support one another, and we cannot ask whether they support one another more in one situation than in the other without having a particular measure of support in mind. And the problem is that the measures of confirmation we considered give very different outcomes here – which is precisely why C_d , C_r , and C_l disagree so starkly on this case. Thus, we must revert to our more general intuitions about coherence. And from that perspective, I think there is an intuitive pull toward saying that in the first situation S is much less coherent than in the second. Although it is tricky to put down your finger on what precisely motivates this intuition, it seems that here we are influenced at least partly by considerations of relative overlap. In the second situation, the overlap of the set of Japanese suspects and that of suspects owning a Samurai sword is very large: there are 9 Japanese owners of a Samurai sword versus only one non-Japanese owner of such a sword and one Japanese non-owner. In the first situation, however, the relative overlap is extremely small: while there are still 9 suspects that both are Japanese and own a Samurai sword, there are 1,050 non-Japanese owners of a Samurai sword and just as many Japanese non-owners of a Samurai sword.

It may seem awkward to introduce considerations of relative overlap in an analysis of coherence as mutual support, especially because the next chapter will analyze coherence as relative overlap. However, since we are now concerned with our general intuitions with respect to coherence, nothing obstructs our use of other considerations than those inspired by a view on coherence as mutual support. Nonetheless, precisely because this analysis does not aim at capturing *all* of our relative-overlap intuitions, I need not endorse the general claim that such considerations should always play a vital role in deciding intuitively which set is more coherent, even if only in cases where the measures of confirmation fail to provide a unanimous answer to the question whether the given propositions support one another more in one situation than in another. Instead, different considerations may influence our intuitive judgements in different cases.

As an illustration of this, consider a third situation in which we have 12 suspects who all live in the same house, and that 10 of them are Japanese, 10 own

a Samurai sword, and 9 are both Japanese and Samurai sword owners. One can verify that here, too, the measures of confirmation do not answer unanimously the question whether J and O support one another more or less in this third situation than they do in situation I of the example above. So, if relative-overlap considerations were to be the arbiter in this kind of case, then, since the relative overlap of J and O is the same in situation III as it is in situation II, we would be compelled to say that in situation III, $\{J, O\}$ is more coherent than it is in situation I. Yet here I believe we do not feel that relative-overlap considerations have the same force in the new example as they have in the example in the text. Rather, our intuition says that both in situation I and in situation III, $\{J, O\}$ is not very coherent and that it is not crystal-clear which of situations I and III makes $\{J, O\}$ more coherent. The intuition that in situation III, $\{J, O\}$ is not very coherent may be due to the fact that in situation III, J and O are almost independent; if, for example, the one person that is now assumed to be a non-Japanese non-owner of a Samurai sword were to own a Samurai sword as well, then the propositions would even be negatively related.¹³

But regardless of whether considerations of overlap are the correct explanation of our intuition with respect to situations I and II, I think that the Samurai example sketched above constitutes a counterexample against measures C_r and C_l and thus in favor of measure C_d . Let me at once admit that I do not feel that this intuition is strong enough to give us *strong* reason to prefer C_d . However, it does seem to tip the balance somewhat in favor of the former measure. If we were to find an example in which C_d behaves counterintuitively, while one of the other two measures does not, it may turn out to be the case that one of the other two options is the most satisfactory after all. Naturally, in that case, we should first study the other measures of confirmation and see what the results would be of using them for our measure of coherence.

As a last remark, let me note a striking asymmetry between the measures of coherence as mutual support presented above. It can easily be checked that C_r and C_l both have range $[0, \infty)$. This may be unsurprising, since r and l have likewise range $[0, \infty)$. But whereas d has range $(-1, 1)$, it can be verified that C_d has range $(-\frac{1}{2}, 1)$.¹⁴ Nonetheless, this need not worry us too much, since there exist ordinally equivalent measures for each of the measures r , l and d which lead to all of the measures C_m having range $(-1, 1)$. These ordinal equivalents will be discussed at the end of the next chapter.

¹³I am much indebted to Mark Siebel for pressing Igor Douven and me to be more explicit about what role is to be given to relative-overlap considerations in this sort of examples.

¹⁴In short, this is due to the fact that if $d(\langle S, S' \rangle)$ approximates -1 , it follows that $d(\langle S', S \rangle)$ approximates 0 . More generally, for any S and S' , the average of $d(\langle S, S' \rangle)$ and $d(\langle S', S \rangle)$ cannot be equal to or lower than $-1/2$ and thus the lower bound of C_d - which takes the average of all such pairs - is likewise $-1/2$. On the other hand, when $d(\langle S, S' \rangle)$ approximates 1 , it can still be the case that $d(\langle S', S \rangle)$ approximates 1 also.

2.5 Comparing C_d With the Extant Measures of Coherence

In this chapter I explicated coherence in the sense of mutual support and presented three potential measures of coherence, one of which appeared to be the most satisfactory. However, several other probabilistic accounts of coherence have been proposed in the literature. Although none of these have been explicit attempts to explicate coherence as mutual support,¹⁵ this does not mean that they cannot capture this sense of coherence, or cannot outperform my proposal C_d . The next four subsections will consider each of the four proposals to be found in the literature.

2.5.1 Shogenji's Measure S

Shogenji's (1999) was the first to propose a probabilistic measure of coherence. His measure was constructed as a response to the argument by Peter Klein and Ted Warfield (1994) and (1996) that coherence cannot be truth conducive *per se*. Clearly, then, Shogenji's proposal seems to be a proposal for a measure of coherence as truth conduciveness much more than for coherence as mutual support. Nonetheless, this should not directly discredit his measure as a measure of coherence as mutual support.

For a set $S = \{R_1, \dots, R_n\}$, Shogenji's probabilistic measure of coherence is defined thus:

$$S(S) =_{df} \frac{p(R_1 \wedge \dots \wedge R_n)}{\prod_{i=1}^n p(R_i)}. \quad (2.7)$$

That is, the coherence of a set is the probability that all of its propositions are true divided by the product of the marginal probabilities of all of its propositions. It is easy to see that for sets consisting of two propositions, this measure is actually equivalent to measure C_r . For a set $S = \{R_1, R_2\}$, $C_r(S)$ gives:

$$C_r(S) = \frac{1}{2} \left(\frac{p(R_1 | R_2)}{p(R_1)} + \frac{p(R_2 | R_1)}{p(R_2)} \right) \quad (2.8)$$

$$= \frac{p(R_1 \wedge R_2)}{p(R_1)p(R_2)} \quad (2.9)$$

$$= S(S) \quad (2.10)$$

However, for three and more propositions this is no longer the case. Instead, Shogenji's measure is not sensitive to relations of the any-any coherence type. Therefore, if one agrees with my arguments for choosing a measure from class ϵ

¹⁵Although Fitelson has proposed to analyze coherence as mutual support, his object is to explicate coherence as the generalization of logical equivalence. More on this in the next chapter.

above, then Shogenji's measure is plainly unsatisfactory as a measure of coherence as mutual support.

2.5.2 Olsson's Measure \mathcal{O}

Erik Olsson (2002) is equally concerned with Klein and Warfield's argument that coherence cannot be truth conducive. However, contrary to Shogenji, Olsson believes that Klein and Warfield are right and that coherence cannot be truth conducive. The question of truth conduciveness will be pivotal in chapter 4 of this thesis, but Olsson also proposes a measure of coherence and, therefore, we have another competitor for a measure of coherence as mutual support.

For a set $S = \{R_1, \dots, R_n\}$, Olsson's (2002: 249) measure reads as follows:

$$\mathcal{O}(S) =_{df} \frac{p(R_1 \wedge \dots \wedge R_n)}{p(R_1 \vee \dots \vee R_n)}.$$

That is, the coherence of a set is the probability that all of its propositions are true divided by the probability that at least one of its propositions is true. Or, in set-theoretical terms, it is the ratio of the area where all the propositions overlap and the total area encompassed by all the propositions. The latter definition has the advantage that it makes clear what Olsson's measure does: it measures the relative overlap of the propositions in a set. As such it can never be a mutual support measure: there will be sets in which all the propositions are independent or even negatively dependent that are more coherent as measured by \mathcal{O} than sets in which all of the propositions are mutually supportive.

To illustrate this, consider the following example of an independent set, originally due to Bonjour (1985: 96) and discussed in Bovens and Hartmann (2003a: 40-43). Consider set $S = \{C, E, T\}$ with

- C : This chair is brown;
- E : Electrons are negatively charged;
- T : Today is Thursday.

Assume that the propositions in this set are independent. This squares well with intuition: the fact that it is Thursday does not seem to affect the color of the chair or the charge of electrons, nor vice versa. In the original example, Bonjour compares this set with a set in which some of the propositions seem, intuitively, to support each other. Thus consider the set $S' = \{A, R, B\}$ with

- A : All ravens are black;
- R : This bird is a raven;
- B : This bird is black.

According to Bonjour (1985: 96), set S possesses 'a very low degree of coherence,' while in set S'

the component propositions, rather than being irrelevant to each other, fit together and reinforce each other in a significant way; from an epistemic standpoint, any two of them would lend a degree of positive support to the third (though only very weak support in two out of the three cases). (*ibid.*)

Bovens and Hartmann (2003a) agree with this. According to them '[t]here is no doubt that set $[S']$ is more coherent than set $[S]$ ' (*op. cit.*: 29). Both Bonjour and Bovens and Hartmann present their conclusion without any qualifications. Apparently, as long as the propositions in the first set are independent and as long as any two propositions in the second set support the third, the second set should be more coherent than the first. More specifically, neither Bonjour nor Bovens and Hartmann let their conclusions depend on the marginal probabilities of the propositions in these sets: as long as they are independent in the first set and positively dependent in the second, it seems to matter little how probable they are.¹⁶

Nonetheless, the problem with this example is that the second set still consists of two pairwise independent propositions and, consequently, that the set is not any-any coherent as defined by definition 2.4. Moreover, as the quote above by Bonjour already indicates, the mutual support relations do not seem to be very strong in two of the three cases. Therefore consider a similar set $S'' = \{B, O, M\}$ with

- B : This bird is black;
- O : This bird is a crow;
- M : This bird has a lifelong mate.

Presumably, in this case it can be assumed that each subset of propositions greatly supports any of the other subsets. More precisely, assume that in the population we are studying most of the black birds are crows and have lifelong mates, most crows are black and have lifelong mates and most birds that have lifelong mates are crows. For all probability models in which this is the case, Bonjour's example is even strengthened: we now have a set in which each proposition is supported *both* by each of the other propositions *and* by the combination of them. And consequently, from the standpoint of coherence as mutual support, it should be the case that the second set is more coherent than the first for all probability models such that the second is an any-any coherent set and the first is an independent set.

¹⁶Although Bovens and Hartmann present their conclusion in an unqualified manner, in the probability model that they later construct to represent these two sets, all propositions have almost the same probability. Therefore, they may have intended their conclusion to hold only in a *ceteris paribus* way, i.e., for instance, given (almost) equal marginal probabilities. Nevertheless, it will appear below that for their own theory, the conclusion does not hold for any of the most plausible *ceteris paribus* conditions.

However, on Olsson's measure this is not the case. For suppose that the probability distributions of sets S and S'' are as in figure 2.4. It can easily be




Figure 2.4: Diagrams of the probability distributions corresponding to sets S (left) and S'' (right)

checked that for these probabilities set S is independent in the sense defined above, while set S'' is any-any coherent. However, on Olsson's measure set S is much more coherent than set S'' : $\mathcal{O}(S) \approx .516$ and $\mathcal{O}(S'') \approx .113$. Clearly, from the view of coherence as mutual support this is not satisfactory.

Let me note in passing that Olsson's measure does satisfy the requirement that any-any coherent sets are more coherent than independent sets in a *ceteris paribus*-sense. For if two sets have an equal number of propositions and if these propositions have equal marginal probabilities, then it will always be the case that an any-any coherent set has a higher degree of coherence as measured by Olsson's measure than an independent set.

For consider a set $S = \{R_1, \dots, R_n\}$ that is any-any coherent and a set $S' = \{R'_1, \dots, R'_n\}$ that is independent and assume that all propositions have the same probability in both sets. In that case, $p(R_1 \wedge \dots \wedge R_n) > p(R'_1 \wedge \dots \wedge R'_n)$ and $p(R_1 \vee \dots \vee R_n) < p(R'_1 \vee \dots \vee R'_n)$, from which it follows directly that $\mathcal{O}(S) > \mathcal{O}(S')$.

2.5.3 Fitelson's Measure C_k

Of all the measures of coherence that have been proposed in the literature so far, Fitelson's measure of coherence is the only measure that fits the general framework presented in section 2.3 above. That is, his measure C_k also measures

coherence as the average of the degree of any-any coherence of a set.¹⁷ However, instead of using one of the three measures of confirmation discussed above, Fitelson proposes to measure support by means of the Kemeny–Oppenheim measure k :¹⁸

$$k(\langle S, S' \rangle) =_{df} \frac{p(S' | S) - p(S' | \neg S)}{p(S' | S) + p(S' | \neg S)}. \quad (2.11)$$

Clearly, the measure that results when we substitute k in definition 2.7 is a measure of mutual support. Furthermore, as the following theorem shows, the measures k and l are ordinally equivalent, where two generalized measures of confirmation m and m' are ordinally equivalent if for all subsets S' and S^* of a set of propositions $S = \{R_1, \dots, R_n\}$ and subsets T' and T^* of a set of propositions $T = \{Q_1, \dots, Q_n\}$, $m(\langle S', S^* \rangle) > m(\langle T', T^* \rangle)$ iff $m'(\langle S', S^* \rangle) > m'(\langle T', T^* \rangle)$ (for a proof see Appendix 2 C, for further discussion see also Fitelson (2005a)):

Theorem 2.3 *The measures k and l are ordinally equivalent for all pairs of sets of propositions $\langle S', S^* \rangle$ and $\langle T', T^* \rangle$ such that $p(S' | \neg S^*) > 0$ and $p(T' | \neg T^*) > 0$.*¹⁹

This would suggest that since we only found a reason to weakly prefer C_d over C_l , we would similarly only have a reason to weakly prefer C_d over C_k . However, the fact that k and l are ordinally equivalent does not guarantee that C_d and C_l are also ordinally equivalent as defined in section 2.3. Instead, as the following theorem shows, this is not the case (for a proof see Appendix 2 C):

Theorem 2.4 *The measures of coherence C_k and C_l are not ordinally equivalent.*

Moreover, I believe we have a good reason to strongly prefer C_l over C_k . For consider again the Samurai example above. I argued there that situation I is intuitively less coherent than situation II. It appeared that on the likelihood measure both situations were almost equally coherent. This holds true for Fitelson's measure C_k as well. However, in both situations it comes out as being close to the maximal degree of coherence a set can have on Fitelson's measure (which is 1, as can easily be ascertained): for situation I, $C_k(S) \approx .97559$ and for situation II, $C_k(S) \approx .97561$. This strikes me as being wrong in at least the first situation,

¹⁷Strictly speaking, this is not true for his (2003) definition, but only for his (2004) definition of his measure C_k . In what follows I will always consider the latest version of his measure as described in his (2004).

¹⁸In fact this is only the definition for the cases in which S contains contingent propositions only. However, the clauses for the other cases can be ignored for the purposes of this chapter. I will present the complete definition in the next chapter (section 3.2, definition 3.1).

¹⁹The reason for the additional constraints that $p(S' | \neg S^*) > 0$ and that $p(T' | \neg T^*) > 0$ is the fact that the likelihood measure is not defined for cases in which these constraints do not hold. See also note 12 above.

given that there the relative overlap between the set of Japanese suspects and the set of suspects owning a Samurai sword is very small, so that, intuitively, one would suppose it to be extremely easy to *considerably* increase the degree of coherence of S – namely, simply by increasing the amount of overlap of the two sets of suspects. But, of course, on Fitelson’s measure there is only room for a very insubstantial increase. This gives a strong reason to prefer measure C_l over C_k , for in the former there still is much room left for an increase in coherence.

2.5.4 Bovens and Hartmann’s Difference Function f_r

Finally, there is Bovens and Hartmann’s (2003a) and (2003b) account of coherence. Just as the other proposals, Bovens and Hartmann’s measure is not intended to be a measure of coherence as mutual support. Instead, Bovens and Hartmann believe that our intuitive judgements of coherence ‘rest on the subtle interplay between the degree of positive relevance relations and relative overlap relations between propositions’ (Bovens and Hartmann 2003a: 53). Moreover, they are concerned with coherence as a confidence boosting property which introduces a number of difficulties which Bovens and Hartmann solve in a very ingenious manner. One of the differences between their approach and the one followed in this chapter is that instead of defining their measure for a set of propositions, they make the additional assumption that one has been informed about each proposition in any given set by a separate, partially reliable witness.

Nonetheless, just as the above measures, their measure may indeed do justice to our mutual-support intuitions even if it has not specifically been designed to do so. Therefore, let us see how satisfactory their theory of coherence can explicate the concept of coherence in the sense of mutual support and postpone a more elaborate discussion of their theory to chapter 4.

To this effect, assume that we are informed about each proposition in any given set by a separate witness. These witnesses are assumed to be independent of one another and to be equally reliable to a degree $r \in (0, 1)$, where $r = 1$ indicates full reliability and $r = 0$ full unreliability. Next, define the relation of *being no less coherent than*, ‘ \succcurlyeq .’ Let ‘ a_i ’ be the sum of the probabilities of all conjunctions of $n - i$ elements of S and the negations of the remaining i elements of S . For a set $S = \{R_1, \dots, R_n\}$, the function $c_r(S)$ is given by:

$$c_r(S) = \frac{a_0 + (1 - a_0)(1 - r)^m}{\sum_{i=0}^m a_i (1 - r)^i}. \quad (2.12)$$

Then for any two sets $S = \{R_1, \dots, R_m\}$, $S' = \{R'_1, \dots, R'_n\}$, the relation ‘ \succcurlyeq ’ is defined as follows:

$$S \succcurlyeq S' \quad \text{iff} \quad f_r(S, S') \geq 0 \text{ for all values of } r \in (0, 1), \quad (2.13)$$

with $f_r(S, S')$ – called the *difference function* – being the difference between the values of c_r for the respective sets:

$$f_r(S, S') = c_r(S) - c_r(S') \quad (2.14)$$

$$= \frac{a_0 + (1 - a_0)(1 - r)^m}{\sum_{i=0}^m a_i(1 - r)^i} - \frac{a'_0 + (1 - a'_0)(1 - r)^n}{\sum_{i=0}^n a'_i(1 - r)^i} \quad (2.15)$$

(a'_0 indicates the probability of the conjunction of the elements of S' , etc.) As Bovens and Hartmann note, ‘ \succsim ’ merely induces a so-called *quasi-ordering* on the set of information sets, that is, an ordering that is reflexive and transitive but not necessarily complete. This means that there may be sets S and S' such that neither $S \succsim S'$ nor $S' \succsim S$. According to them this fact does justice to our intuitions: intuitively, we are sometimes unable to decide which of two sets is the more coherent, even in cases in which it is not clear either that they have equal coherence.²⁰

For example, in situations II and III of the Samurai example above (section 2.4) it is not clear at all that the set has an equal degree of coherence in both situations: they are very different from each other, so we have little reason to expect the degree of coherence in both cases to be precisely the same. But if any of the two sets is more coherent than the other, intuitively, it seems hard to say which one it is. Measure C_d solves this problem by ruling that indeed one of the two situations is more coherent than the other, but that the difference is minimal ($\approx .06$ on a $(-1/2, 1)$ scale). Arguably, this is a satisfactory explanation for why our intuitions are at a loss in these situations. However, an approach that leaves such cases indeterminate seems to be equally satisfactory.

Nonetheless, this does not mean that it would never count against their theory if it produced an ‘indeterminate’ judgement in a specific example. For example, from the standpoint of coherence as mutual support, an any-any coherent set should always be more coherent than an independent set, even if the difference is but small. Bovens and Hartmann may be expected to agree, since, as discussed above, they believe that in Bonjour’s example the independent set is very clearly less coherent than the set in which some of the propositions support each other and none of the propositions undermine each other.

However, from the example used to discredit Olsson’s theory as a theory of coherence as mutual support it follows that this conclusion does not hold for Bovens and Hartmann’s difference function either. From the values in diagram 2.4, it follows that in this case, for set S , we have $a_0 = .512$, $a_1 = .384$, $a_2 = .096$ and $a_3 = .008$, and for S'' , we have $a''_0 = .016$, $a''_1 = .126$, $a''_2 = 0$ and $a''_3 = .858$. So

²⁰Bovens and Hartmann (2003a: 40) present one example in which we are, according to them, unable to decide which of two sets is the more coherent. I will discuss this example briefly in section 4.6.

the difference function for the two sets is this:

$$f_r(S, S'') = \frac{.512 + .488(1-r)^3}{.512 + .384(1-r) + .096(1-r)^2 + .008(1-r)^3} - \frac{.016 + .984(1-r)^3}{.016 + .126(1-r) + .858(1-r)^3} \quad (2.16)$$

And, as figure 2.5 shows, the graph of this function crosses the r -axis. Therefore, it is not the case that $f_r(S, S'') > 0$ for all values of $r \in (0, 1)$ and so neither of the two sets is more coherent than the other according to Bovens and Hartmann's theory of coherence.




Figure 2.5: $f_r(S, S'')$

However, the case is even worse: while for Olsson's measure the requirement that any-any coherent sets are more coherent than independent sets remains true in a *ceteris paribus* sense, for Bovens and Hartmann not even this is the case. For consider the new probability distribution for sets $S = \{C, E, T\}$ and $S'' = \{B, O, M\}$ (as defined above) in figure 2.6. It can easily be verified that in this case the marginal probabilities of all propositions are equal: all propositions have a marginal probability of .1. However, as can be seen from figure 2.7, the case is still indeterminate.

Here one may feel that different *ceteris paribus* conditions might work. For example, one may propose that instead of the unconditional probabilities, the probabilities that all propositions are true should be equal, i.e., $a_0 = a_0''$. However, in that case similar examples are possible. For example, consider any probability model for set S'' in which $a_0'' = .001$, $a_1'' = .03$ and $a_2'' = 0$ and compare this with the independent set S above. In this case, $a_0 = a_0''$, while S'' is still any-any coherent. Still, $f_r(S'', S) < 0$ for $r > .9875$.

Alternatively, one may propose that the value of a_n must remain equal, which means, informally put, that the total area within the circles of the propositions




Figure 2.6: New diagrams of the probability distributions corresponding to sets *S* (left) and *S'* (right)

must remain equal. I cannot see any justification for that move, although in Bovens and Hartmann's discussion of Bonjour's example the values for a_n and a'_n are indeed equal. However, this alternative would not work either. For in this case take any probability model for set *S'* in which $a''_0 = .122$, $a''_1 = .149$ and $a''_2 = 0$ and compare this set with *S*. It can easily be checked that *S'* is any-any coherent and that $a''_n = a_n$. However, $f_r(S'', S) < 0$ for $r < .14$.




Figure 2.7: $f_r(S'', S)$ for the new probability distribution

2.6 Anticipated Objections

Above (section 2.5.3) I noted that Fitelson's approach is the only approach to measuring coherence that can be made to fit into the general framework presented in sections 2.2 and 2.3. In this section I will consider two different arguments which have been leveled against Fitelson's measure. It will appear that both arguments are not specific to Fitelson's measure but apply to all measures that can be formulated using the above scheme. Nevertheless, I will show that in both cases the intuitions underlying the arguments are not valid within an explication of coherence as mutual support. To forestall misunderstanding, let me emphasize that these are not the only objections that have been or can be raised against measures of coherence as mutual support. In chapter 6, I will consider some objections that could be leveled against all of the measures of coherence proposed in this thesis, and, consequently, also against measures of coherence as mutual support.

2.6.1 Siebel's Objection Against Fitelson's Measure

Consider the following example by Mark Siebel (2004). Suppose that there are ten suspects in a murder case. Suppose furthermore that each of them has a .1 probability of having committed the crime and that each of them has committed at least one earlier crime: eight of them have committed pickpocketing, eight have committed a robbery, and six have done both. Now consider the set $S = \{P, R\}$ with

- P : The murderer has committed pickpocketing;
- R : The murderer has committed a robbery.

The probability model for this set is given by figure 2.8. According to Siebel (2004: 190), the mere fact that there is 'a strong coincidence' between the two propositions makes set S intuitively coherent. Presumably, the strong coincidence between the propositions is due to the fact that most of those who have committed pickpocketing have also committed a robbery, and vice versa. And since on Fitelson's measure S comes out as being incoherent ($C_k \approx -0.14$), Siebel believes we thus have a counterexample against Fitelson's measure. Since all of the above measures of coherence as mutual support agree with Fitelson's measure that the set is incoherent, we would have a counterexample against all measures of coherence as mutual support.

I indicated above that when we explicate a specific sense of coherence we must be careful in applying our basic intuitions with respect to the general concept of coherence. Although in some cases it seems that such intuitions may provide a tool for distinguishing the most satisfactory measures, in other cases it may be





Figure 2.8: Diagram of the probability distribution for set S

that the intuitions are fueled by one of the alternative senses of coherence. For instance, the example with the independent and the any-any coherent sets by Bonjour may be an example which is at home within the concept of coherence as mutual support, but not in the sense of coherence as relative overlap. As the counterexample against Olsson's measure shows, the relative overlap of the propositions of an independent set may be much larger than that of an any-any coherent set. Alternatively, some intuitions may be informed by the sense of coherence as relative overlap and not be applicable to coherence in the sense of mutual support. This, I believe, is what is going on in this example.

Presumably, the two propositions in set S 'coincide' because they overlap for a large part. However, this does not directly imply that the set is any-any coherent as defined by definition 2.4. If in a set both propositions have a $.9$ marginal probability, then they still have a $.81$ overlap if they are independent. If they have a $.99$ probability, the overlap in case of independence is even $.9801$ and, therefore, an overlap of less than $.98$ would imply that the propositions, though overlapping very much, actually undermine each other. And this is precisely what is going on in Siebel's example.

As Siebel (2004: 190) notes, the set $S = \{P, R\}$ is an example of a subcontrary set, which is a set in which all propositions cannot be jointly false. A quick calculation shows that subcontrary sets consisting of two propositions are always any-any incoherent. Consequently, a measure of coherence as relative overlap should be negative for the above example: each any-any incoherent set should have negative coherence. Therefore, Siebel's intuition that the set in his example

Figure 2.9: $c_r(S)$ for values of $.01 \leq x \leq .9$

is coherent cannot follow from a view of coherence as mutual support. Instead, his conviction seems to be based on the intuition that a large overlap merits a positive degree of coherence, an intuition that appeals to a conception of coherence as relative overlap.

2.6.2 Bovens and Hartmann's Objection Against Fitelson's Measure

Bovens and Hartmann (2003a: 50 f) consider a set consisting of two propositions, R_1 and R_2 , for which they take the probability of their conjunction – their ‘overlap’ in set-theoretic terms – variable, $p(R_1 \wedge R_2) = x$, and set both $p(R_1 \wedge \neg R_2)$ and $p(\neg R_1 \wedge R_2)$ equal to .05. According to them, increasing x from .01 to .8 should monotonically increase the coherence of the set: ‘[i]ntuitively, one would think that when keeping the non-overlapping area fixed, then, the more overlap, the greater the coherence’ (Bovens and Hartmann 2003a: 51). Figure 2.9 shows the graph of $c_r(S)$ for values of x between .01 and .9. Evidently, the graph increases monotonically and therefore, the set will be more coherent in case x is larger.

As Bovens and Hartmann note, on Fitelson’s measure this is not the case:

[t]he measure first increases from $[p(R_1 \wedge R_2) = .01]$ and then reaches its maximum for $[p(R_1 \wedge R_2) \approx .17]$ and subsequently decreases again. We fail to see any intuitive justification for this behaviour of the measure. (Bovens and Hartmann 2003a: 51)

It can easily be checked that precisely the same objection holds for C_d , which likewise increases until it reaches its maximum for $x \approx .17$ and decreases afterwards. The graphs of these measures are given in figure 2.10.

What happens in the example is that the probability that both propositions are true increases, while the probability that precisely one of the two proposi-




Figure 2.10: C_k (left) and C_d (right) for values of $.01 \leq x \leq .9$

tions is false remains the same. In some cases this strategy intuitively increases the coherence of a set, most notably if two inconsistent propositions are made consistent by increasing the value of x from 0 to a value $x_0 > 0$. Nevertheless, one should also expect coherence to decrease for certain values of x , precisely because for $x = .9$ the set in this example has become a subcontrary set. Since, as noted above, the coherence of a subcontrary set consisting of two propositions must be negative, one would expect the coherence to decrease as x approaches .9. Thus, from the point of view of coherence as mutual support, coherence should not be a monotonically increasing function of x .

The replies to Siebel's (2004) and Bovens and Hartmann's (2003a) examples do not discredit the intuitions that they appeal to *per se*. However, they do show that these intuitions are not at home in an analysis of coherence as mutual support. Instead, it will appear that a measure of coherence as relative overlap does in fact do justice to those intuitions. Whether or not it is possible to construct an account of coherence that does justice to both intuitions of mutual support and relative overlap will be considered in sections 3.3, 3.6 and 3.7.

2.7 Conclusion

In this chapter I have explicated the notion of coherence as mutual support. From the general discussion in the first sections it appeared that there are a large number of candidates for the explication of coherence as mutual support. Therefore, I have limited my discussion in a number of ways. Firstly, I have argued that one should consider only those measures that measure the degree of any-any coherence. Secondly, I have preferred measures that take a straight average over measures that take a weighted average of the different elements. And thirdly, I have only discussed measures of coherence that are based on the three most popular measures of confirmation. I believe I have given good reasons for the first limitation, but not for the other two. Most notably, it remains to be seen that

none of the other measures of confirmation performs better than the measures of coherence discussed in this chapter.

In order to decide which of the measures of confirmation provides the best basis for a measure of coherence, the analysis of the first sections of this chapter will not be of much help. That is, for all measures of confirmation c such that $c(H, H')$ is positive if H probabilistically confirms H' , negative when H probabilistically disconfirms H' and zero when the two are independent, the general discussion above will not tell us which of the measures gives the most satisfactory analysis of coherence.

In order to bring the number of potential measures of coherence down to a single measure, we need to bring in our basic intuitions with respect to coherence. On the basis of an example purportedly displaying one of these intuitions, I have made a tentative case for measure C_d .

Also, I have considered the extant accounts of coherence and discussed in how far they are satisfactory as measures of coherence as mutual support. It appeared that none of the accounts constitutes a satisfactory explication of this sense of coherence. Shogenji's measure is not sensitive to coherence relations of the any-any type, while Olsson's and Bovens and Hartmann's measures are not measures of mutual support at all. Finally, Fitelson's account, while definitely a measure of mutual support, also has a strikingly unsatisfactory consequence. Therefore, the tentative conclusion of this chapter must be that C_d provides the most satisfactory explication of coherence in the sense of mutual support.

Finally, I have considered two objections that have been proposed against Fitelson's measure of coherence and that are equally applicable to the measures of coherence as mutual support proposed in this chapter. While it appeared that these objections are not valid from the point of view of coherence as mutual support, they do seem to suggest that the concept of coherence as mutual support does not respect all of our intuitive judgements of coherence. In the next chapter, I will propose a measure of coherence as relative overlap that respects the intuitions underlying the examples put forward by Bovens, Hartmann and Siebel.

Appendix 2 A: Proof of Theorem 2.1

Theorem 2.1

1. *One-all coherence does not imply one-any coherence.*
2. *One-all coherence does not imply partition coherence.*
3. *One-any coherence does not imply any-any coherence.*
4. *One-any coherence does not imply partition coherence.*
5. *Partition coherence does not imply any-any coherence.*
6. *Partition coherence does not imply one-any coherence.*

Proof: Theorem 2.1.1 follows from the first probability model, theorems 2.1.2, 2.1.3, and 2.1.4 follow from the second probability model, and theorems 2.1.5 and 2.1.6 follow from the third.

Model 1:

R_1	R_2	R_3	probability	R_1	R_2	R_3	probability
T	T	T	.00296053	F	T	T	.000103878
T	T	F	.00131579	F	T	F	.0811461
T	F	T	.000986842	F	F	T	.009375
T	F	F	.894737	F	F	F	.009375

Model 2:

R_1	R_2	R_3	R_4	probability	R_1	R_2	R_3	R_4	probability
T	T	T	T	.000938319	F	T	T	T	.0204122
T	T	T	F	.00404002	F	T	T	F	.0058595
T	T	F	T	.00542335	F	T	F	T	.00447617
T	T	F	F	.0208483	F	T	F	F	.0630022
T	F	T	T	.00539031	F	F	T	T	.0159602
T	F	T	F	.0225303	F	F	T	F	.0498692
T	F	F	T	.0289642	F	F	F	T	.0434353
T	F	F	F	.0368652	F	F	F	F	.671985

Model 3:

R_1	R_2	R_3	R_4	probability	R_1	R_2	R_3	R_4	probability
T	T	T	T	.000732422	F	T	T	T	.0000610352
T	T	T	F	.00268555	F	T	T	F	.00335693
T	T	F	T	.000549316	F	T	F	T	.0133057
T	T	F	F	.0272827	F	T	F	F	.0770264
T	F	T	T	.00341797	F	F	T	T	.00964355
T	F	T	F	.0244141	F	F	T	F	.0806885
T	F	F	T	.0265503	F	F	F	T	.0707397
T	F	F	F	.0393677	F	F	F	F	.620178

■

Appendix 2 B: Proof of Theorem 2.2

Theorem 2.2 *For any pair of distinct $m, m' \in \{d, r, l\}$ there exist infinitely many triples $\langle S, S', p(\cdot) \rangle$ such that either $C_m(S) > C_m(S')$ and $C_{m'}(S) < C_{m'}(S')$ or $C_m(S) < C_m(S')$ and $C_{m'}(S) > C_{m'}(S')$.*

Proof: Consider the class of all sets containing exactly three propositions for which the following hold (where I use r_1, r_2 , and r_3 as propositional variables):

- $p(r_1 \wedge r_2 \wedge r_3) \in (0, .1875)$;
- $p(r_1 \wedge r_2 \wedge \neg r_3) = p(r_1 \wedge \neg r_2 \wedge r_3) = p(\neg r_1 \wedge r_2 \wedge r_3) = p(r_1 \wedge r_2 \wedge r_3) + 1/20$;
- $p(r_1 \wedge \neg r_2 \wedge \neg r_3) = p(\neg r_1 \wedge r_2 \wedge \neg r_3) = p(\neg r_1 \wedge \neg r_2 \wedge r_3) = 1/30$.

Without loss of generality, assume that for every $x \in (0, .1875)$ there is exactly one set in the just-defined class such that the probability of the conjunction of its members equals x , and denote that set by $S(x)$. We then can write the formulas for the three measures as functions of x , as follows:

$$\begin{aligned}
 C_d(S(x)) &= \frac{3}{12} \left(-\frac{2}{15} - 3x + \frac{x}{\frac{1}{20} + 2x} \right) + \frac{3}{12} \left(-\frac{1}{20} - 2x + \frac{x}{\frac{2}{15} + 3x} \right) \\
 &\quad + \frac{6}{12} \left(-\frac{2}{15} - 3x + \frac{\frac{1}{20} + 2x}{\frac{2}{15} + 3x} \right) \\
 C_r(S(x)) &= \frac{6}{12} \left(\frac{\frac{1}{20} + 2x}{\left(\frac{2}{15} + 3x\right)^2} \right) + \frac{6}{12} \left(\frac{x}{\left(\frac{1}{20} + 2x\right)\left(\frac{2}{15} + 3x\right)} \right) \\
 C_l(S(x)) &= \frac{3}{12} \left(\frac{\left(\frac{19}{20} - 2x\right)x}{\left(\frac{1}{20} + 2x\right)\left(\frac{2}{15} + 2x\right)} \right) + \frac{3}{12} \left(\frac{\left(\frac{13}{15} - 3x\right)x}{\left(\frac{1}{20} + x\right)\left(\frac{2}{15} + 3x\right)} \right) \\
 &\quad + \frac{6}{12} \left(\frac{\left(\frac{13}{15} - 3x\right)\left(\frac{1}{20} + 2x\right)}{\left(\frac{1}{12} + x\right)\left(\frac{2}{15} + 3x\right)} \right)
 \end{aligned}$$

One easily verifies that for x going from 0 to .1875, $C_d(S(x))$ first monotonically increases until it reaches its maximum at $x \approx .0314$ and then monotonically decreases; $C_r(S(x))$ monotonically increases until it reaches its maximum at $x \approx .0166$ and then monotonically decreases; and $C_l(S(x))$ monotonically increases until it reaches its maximum at $x \approx .0177$ and then monotonically decreases. Let $S = \{R_1, R_2, R_3\}$ with $p(R_1 \wedge R_2 \wedge R_3) = a$ and $S' = \{R'_1, R'_2, R'_3\}$ with $p(R'_1 \wedge R'_2 \wedge R'_3) = b$ such that $a, b \in (.0166, .0314)$ and $a < b$; obviously there are infinitely many such pairs of sets. Then it will hold that $C_d(S) > C_d(S')$ but $C_r(S) < C_r(S')$.

In the same way one proves the existence of infinitely many triples $\langle S, S', p(\cdot) \rangle$ for the other combinations of two measures such that these measures disagree in their verdict as to which of S and S' is the more coherent set. ■

Appendix 2 C: Proof of Theorems 2.3 and 2.4

Theorem 2.3 *The measures k and l are ordinally equivalent for all pairs of sets of propositions $\langle S', S^* \rangle$ and $\langle T', T^* \rangle$ such that $p(S' | \neg S^*) > 0$ and $p(T' | \neg T^*) > 0$.*

Proof: Remember that two measures of coherence m and m' are ordinally equivalent if for all subsets S' and S^* of a set of propositions $S = \{R_1, \dots, R_n\}$ and subsets T' and T^* of a set of propositions $T = \{Q_1, \dots, Q_n\}$, $m(\langle S', S^* \rangle) > m(\langle T', T^* \rangle)$ iff $m'(\langle S', S^* \rangle) > m'(\langle T', T^* \rangle)$. Next, rewrite the Kemeny–Oppenheim measure:

$$\begin{aligned} k(\langle S, S' \rangle) &= \frac{p(S' | S) - p(S' | \neg S)}{p(S' | S) + p(S' | \neg S)} \\ &= \frac{\frac{p(S' | S)}{p(S' | \neg S)} - 1}{\frac{p(S' | S)}{p(S' | \neg S)} + 1} \\ &= \frac{l(\langle S, S' \rangle) - 1}{l(\langle S, S' \rangle) + 1}, \end{aligned}$$

in which I have assumed that $p(S' | \neg S) > 0$. Next, consider two subsets S' and S^* of a set of propositions $S = \{R_1, \dots, R_n\}$ and two subsets T' and T^* of a set of propositions $T = \{Q_1, \dots, Q_n\}$. Assume that $p(S' | \neg S^*) > 0$ and $p(T' | \neg T^*) > 0$. Then, one can suppose without loss of generality that $l(\langle S', S^* \rangle) = \alpha$ and $l(\langle T', T^* \rangle) = \beta$. In that case

$$k(\langle S', S^* \rangle) = \frac{\alpha - 1}{\alpha + 1}$$

and

$$k(\langle T', T^* \rangle) = \frac{\beta - 1}{\beta + 1}.$$

Clearly, k and l are ordinally equivalent if

$$\frac{\alpha - 1}{\alpha + 1} > \frac{\beta - 1}{\beta + 1} \quad \text{iff} \quad \alpha > \beta$$

But it is easy to see that this is indeed the case:

$$\begin{aligned} &\frac{\alpha - 1}{\alpha + 1} > \frac{\beta - 1}{\beta + 1} \\ \Leftrightarrow &\frac{\alpha - 1}{\alpha + 1} - \frac{\beta - 1}{\beta + 1} > 0 \\ \Leftrightarrow &\frac{(\alpha - 1)(\beta + 1)}{(\alpha + 1)(\beta + 1)} - \frac{(\alpha + 1)(\beta - 1)}{(\alpha + 1)(\beta + 1)} > 0 \end{aligned}$$

$$\begin{aligned}
&\Leftrightarrow \frac{(\alpha - 1)(\beta + 1) - (\alpha + 1)(\beta - 1)}{(\alpha + 1)(\beta + 1)} > 0 \\
&\Leftrightarrow (\alpha\beta + \alpha - \beta - 1) - (\alpha\beta - \alpha + \beta - 1) > 0 \\
&\Leftrightarrow 2\alpha - 2\beta > 0 \\
&\Leftrightarrow \alpha > \beta
\end{aligned}$$

Therefore, k and l are ordinally equivalent for all pairs of sets of propositions $\langle S', S^* \rangle$ and $\langle T', T^* \rangle$ such that $p(S' | \neg S^*) > 0$ and $p(T' | \neg T^*) > 0$ ■

Theorem 2.4 *The measures of coherence C_k and C_l are not ordinally equivalent.*

Proof: The proof is similar to the proof of theorem 2.2. In this case, consider the class of all sets containing exactly two propositions for which the following hold:

- $p(r_1 \wedge r_2) \in (0, .25)$;
- $p(r_1 \wedge \neg r_2) = .35$;
- $p(r_1 \wedge \neg r_2) = .01$;

By reasoning similar to that in the proof of theorem 2.2, one can verify that for x going from 0 to .25, $C_k(S(x))$ first monotonically increases until it reaches its maximum at $x \approx .148$ and then monotonically decreases and that $C_l(S(x))$ monotonically increases until it reaches its maximum at $x \approx .235$ and then monotonically decreases. Let $S = \{R_1, R_2, R_3\}$ with $p(R_1 \wedge R_2) = a$ and $S' = \{R'_1, R'_2\}$ with $p(R'_1 \wedge R'_2) = b$ such that $a, b \in (.149, .235)$ and $a < b$; again, there exist infinitely many such pairs of sets. Then it will hold that $C_l(S) > C_l(S')$ but $C_k(S) < C_k(S')$. From which it follows directly that C_k and C_l are not ordinally equivalent. ■

Chapter 3

Coherence as Relative Overlap

3.1 Introduction

In the introductory chapter I briefly discussed the view of Kemeny and Oppenheim that the best way to construct a measure of confirmation is by listing enough desiderata to single out one specific class of possible measures. According to them, our intuitions are too vague to ensure that we arrive at the most satisfactory explication of the concept of confirmation and we would be much better off if we arrived at a measure through listing a number of desiderata that such a measure should clearly satisfy. Although their argument seems applicable to measures of coherence as well, I pointed out in section 1.6 that the Kemeny–Oppenheim strategy may be problematic if the general desiderata which are to determine the appropriate measure are justified by an appeal to our intuitions about the concept we are explicating.

For this reason, I dismissed this strategy in the previous chapter and instead chose to compare the different measures of coherence with an example appealing to our basic intuitions. However, some may feel that this dismissal of the Kemeny–Oppenheim strategy was too quick. For there is indeed a general desideratum that has been proposed by both Fitelson (2003) and Bovens and Hartmann (2003a). Indeed, Fitelson uses this desideratum – together with a view that a measure of coherence must be a mutual-support measure – to do something very close to what Kemeny and Oppenheim propose, namely, to bring the class of possible measures of coherence down to one ordinal class. The additional desideratum that Fitelson, Bovens and Hartmann propose says, roughly, that sets consisting of equivalent propositions should have maximal coherence. Although I noted this possibility in section 2.4, I postponed criticizing it until this chapter. The reason for this is that the discussion of this desideratum will make plain that it –

just like Bovens and Hartmann's (2003a) and Siebel's (2004) critiques of Fitelson's measure – is only compatible with a view of coherence as relative overlap.

More specifically, from a discussion of Fitelson's desiderata-based approach, in section 3.2, it will appear in section 3.3 that the combination of the view that sets consisting of equivalent propositions have maximal coherence and the view that a measure of coherence should be a member of class ϵ (as defined in section 2.3) leads to some very counterintuitive consequences. In order to do justice to the former view, we need a measure of coherence as relative overlap, which is defined in section 3.4. For this project the Kemeny–Oppenheim strategy is much more suitable than for measuring coherence as mutual support. It will appear that one single desideratum (which seems very appropriate) is sufficient to determine the appropriate measure up to ordinal equivalence. Moreover, as I will show in section 3.5, the measure proposed in section 3.4 provides a better explication of coherence as relative overlap than any of the extant measures of coherence.

The discussion in this and the previous chapter will leave us with two different measures of coherence as an epistemic virtue. Although some may not find this objectionable – if our intuitions hinge on two different conceptions of coherence, then perhaps we should also have two different explications of coherence – others may agree with Bovens and Hartmann that a measure of coherence should do justice to both of these intuitions at the same time. In that case, we would have one single measure of coherence as an epistemic virtue. The matter of measuring coherence as an epistemic virtue will be the subject of the last two sections of this chapter. First, in section 3.6, I will show that the only measure of coherence that could be interpreted as measuring coherence as both mutual support and relative overlap – Bovens and Hartmann's difference function – is not satisfactory as such. Next, section 3.7 will tentatively propose a specific measure of coherence as an epistemic virtue, which takes a straight average of a measure of coherence as mutual support and a measure of coherence as relative overlap.

3.2 Fitelson's Theory of Coherence

Although Fitelson's measure of coherence can be made to fit the general framework of the previous chapter, his aim is not to explicate our intuitive notion of coherence. Instead, his measure of coherence is intended to be a quantitative generalization of a deductive (logical) concept. In order to make clear what this means it is instructive to first consider Fitelson's views on measures of confirmation.

As is well-known, Carnap (1950) tried to construct a logical theory of confirmation. According to him, '[d]eductive logic may be regarded as the theory of the relation of logical consequence, and inductive logic as the theory of another

concept which is likewise objective and logical, viz., ... degree of confirmation' (Carnap 1950: 43). Confirmation, he argued, can be a logical relation if it is considered to be a quantitative generalization of the logical (or deductive) notion of entailment. In his (2001b) and (2005a), Fitelson combines the Carnapian view of confirmation as generalized logical entailment with the Kemeny–Oppenheim strategy for explicating an intuitive concept (see section 1.6). To this effect, he formulates a number of desiderata for measures of confirmation that he believes flow from our intuitive conception of generalized entailment. Together these desiderata succeed in limiting the number of viable measures to a single ordinal class of measures.

For reasons that will become clear below, it is helpful to divide Fitelson's desiderata into two subsets, which I will call the entailment desiderata and the dependence desiderata. Consider two sets of propositions $S = \{R_1, \dots, R_m\}$ and $S' = \{R'_1, \dots, R'_n\}$ with $m, n > 1$ and a measure of confirmation $c(R, R')$ which gives the degree to which $\bigwedge R'$ confirms $\bigwedge R$, where R and R' are subsets of sets S and S' , respectively. Then the two subsets of desiderata are (cf. Fitelson 2005a):

Entailment Desiderata

(e₁) $c(R, R')$ is maximal and positive if $\bigwedge R'$ entails $\bigwedge R$;¹

(e₂) $c(R, R')$ is minimal and negative if $\bigwedge R'$ entails $\neg \bigwedge R$.

Dependence Desiderata

(d₁) $c(R, R') > 0$ if $p(R|R') > p(R)$;

(d₂) $c(R, R') = 0$ if $p(R|R') = p(R)$;

(d₃) $c(R, R') < 0$ if $p(R|R') < p(R)$.

In Fitelson's view, no other considerations should be taken into account when deciding which measure of confirmation is to be preferred. That is, all measures which satisfy these desiderata are equally suitable. Such an approach would not be very conducive to the project of measuring confirmation if it turned out that a large number of very different measures each satisfied the desiderata. Fortunately, most of the measures of confirmation that can be found in the literature violate the entailment desiderata. For instance, the difference measure, the ratio measure and Carnap's relevance measure (Carnap 1950: §62) are not maximal for equivalent propositions. Furthermore, as Fitelson (2005a) notes, all the

¹Here entailment must be interpreted as logical entailment relative to the background knowledge, i.e., $\bigwedge R'$ entails $\bigwedge R$ iff $K \cup R' \vdash \bigwedge R$, where K is the background knowledge. As in the other chapters, I will suppress reference to the background knowledge.

known measures that *do* satisfy both subsets of desiderata (like, for instance, the Kemeny–Oppenheim measure and the log-likelihood measure) are members of the same ordinal class (as defined in subsection 2.5.3). Thus, given the measures that have been proposed so far, the desiderata are sufficient to determine the desired measure of confirmation up to ordinal equivalence (cf. Fitelson 2005a). From the class of potential measures, Fitelson chooses the Kemeny–Oppenheim (1952) measure of factual support, which has been defined earlier (equation (2.11)). Here is the definition of the Kemeny–Oppenheim measure as a measure of the degree of confirmation of $\bigwedge R$ by $\bigwedge R'$:²

$$F(R, R') =_{df} \frac{p(R' | R) - p(R' | \neg R)}{p(R' | R) + p(R' | \neg R)}. \quad (3.1)$$

It is one question whether there exists a measure satisfying the desiderata; it is another question whether the desiderata are acceptable. To see that this seems indeed to be the case, first note that the members of each of the subsets fit together very naturally. For example, keeping the desideratum that entailment provides maximal confirmation, while rejecting the desideratum that contradiction provides minimal confirmation, makes little sense. And similarly, without the desideratum that independence provides zero confirmation, the desideratum that positive (negative) dependence provides positive (negative) confirmation seems to be ungrounded, and vice versa.

Secondly, the dependence desiderata remain virtually uncontested, since the conditions in (d_1) – (d_3) are generally seen as the standard definitions of probabilistic confirmation, independence, and disconfirmation, respectively. However, the same does not apply to the entailment desiderata above. In an earlier paper with Ellery Eells, Fitelson presented a number of symmetry considerations – all of which seem quite plausible – which may lead to very different measures of confirmation (Eells and Fitelson 2002). So why should we accept Fitelson’s (2005a) preference for the entailment desiderata?

I believe the answer to this question is that the entailment desiderata flow naturally from his view of confirmation as the quantitative generalization of logical entailment. For if confirmation is to be the generalization of logical entailment, then it seems at least initially plausible that logical entailment and logical contradiction should be the extremities of such a measure. Therefore, the dependence and the entailment desiderata together result from his view of confirmation as the *probabilistic* generalization of *logical entailment*.

Apparently, but not surprisingly, the logical notion a certain measure is supposed to generalize plays a crucial role in deciding which measure is appropriate and which is not. However, whereas it is clear which logical notion is generalized

²Again, the definition is for contingent propositions only. The definition of Fitelson’s measure of coherence presented below will apply also to noncontingent propositions.

by confirmation, the case is somewhat more obscure for coherence. In his (2003: 194), Fitelson considered coherence to be the generalization of logical coherence. Unfortunately, of course, there is no generally accepted notion of logical coherence, and therefore it seems hard to see what desiderata would follow from such a concept. In his (2004) he therefore specified his views and suggested that coherence should be taken as the quantitative generalization of logical consistency. At first sight, this seems a natural choice. If one restricts the search for such a notion to the most common logical concepts, then the only alternatives seem to be equivalence and consistency, and it seems that either one of these could be linked to our intuitive concept of coherence. However, a closer look reveals that the concept of consistency has a serious drawback: for unlike equivalence, consistency can be the property both of a theory and of a single proposition, while in general coherence is considered to be a relation *between* propositions and not a property *of* propositions (but see Akiba 2000 and Moretti and Akiba 2005). Therefore, the concept of equivalence seems to be the best choice. Recently, Fitelson has indicated that he, too, has come to hold this view (personal communication).

As observed in the previous chapter, the measure proposed in Fitelson (2004) differs in a number of other ways from his (2003) proposal. In what follows I will discuss the latest version of his theory.

In this proposal, measures of coherence measure the coherence of a set of propositions $S = \{R_1, \dots, R_n\}$. Fitelson's procedure for the construction of his measure consists of two steps: first he defines a measure $c(R, R')$ for the measure in which one subset of propositions $R \subset S$ coheres with another subset of propositions $R' \subset S$ and then he constructs a measure $C(S)$ of the coherence of set S on the basis of these c -values. Clearly, in order to uphold a desiderata-based approach, it is necessary to list desiderata for both c and C . Fitelson's desiderata for c are almost the same as his desiderata for a measure of confirmation.

Equivalence Desiderata

(q_1) $c(R, R')$ is maximal and positive if $\bigwedge R$ and $\bigwedge R'$ are equivalent;³

(q_2) $c(R, R')$ is minimal and negative if $\bigwedge R$ and $\bigwedge R'$ are inconsistent.

Dependence Desiderata

(d_1) $c(R, R') > 0$ if $p(R|R') > p(R)$ (and hence $p(R'|R) > p(R')$);

(d_2) $c(R, R') = 0$ if $p(R|R') = p(R)$ (and hence $p(R'|R) = p(R')$);

(d_3) $c(R, R') < 0$ if $p(R|R') < p(R)$ (and hence $p(R'|R) < p(R')$).

³Where equivalence, too, must be understood as relative to the background knowledge: $\bigwedge R$ and $\bigwedge R'$ are equivalent iff $K \cup R \vdash \bigwedge R'$ and $K \cup R' \vdash \bigwedge R$, where K is the background knowledge.

Clearly, the dependence desiderata remain the same. So the only difference between the desiderata for measures of confirmation and measures of coherence is that for the latter the entailment desiderata are replaced by the equivalence desiderata.

Fitelson's desiderata for C are quite straightforward as well. According to him, a measure of coherence for a set S should simply calculate the average coherence of each subset of S with each disjoint non-empty subset of S (as measured by c). Combining this with the desiderata for c , it follows that a measure of coherence that takes the average of the Kemeny–Oppenheim measures for all disjoint subsets of S satisfies both constraints.

Above, in subsection 2.5.3, I defined Fitelson's measure of coherence as measure C_k , i.e., as the measure which takes the straight average of the Kemeny–Oppenheim measure applied to each of the members of class ϵ . However, the precise definition of Fitelson's measure includes two subclauses for the cases in which at least one of the subsets in set S has either zero or unit probability. Let me present the full definition here. Let $[S]$ again indicate the set of ordered pairs of non-empty non-overlapping subsets of $S = \{R_1, \dots, R_n\}$, that is, $[S] = \{ \langle S', S^* \rangle \mid S', S^* \subset S \setminus \emptyset \wedge S' \cap S^* = \emptyset \}$ and let $[[S]]$ again denote the cardinality of $[S]$. For mnemonic purposes, I will call Fitelson's complete measure $\mathcal{F}(S)$:

Definition 3.1 *Given a set $S = \{R_1, \dots, R_n\}$ and an ordering $\langle \hat{S}_1, \dots, \hat{S}_{[[S]]} \rangle$ of the members of $[S]$, the degree of coherence of S is given by the function*

$$\mathcal{F}(S) =_{df} \frac{\sum_{i=1}^{[[S]]} k(\hat{S}_i)}{[[S]]}, \quad (3.2)$$

with

$$k(S', S^*) =_{df} \begin{cases} \frac{p(S^*|S') - p(S^*|\neg S')}{p(S^*|S') + p(S^*|\neg S')} & \text{if } S^* \not\vdash S' \text{ and } S^* \not\vdash \neg S'; \\ 1 & \text{if } S^* \vdash S' \text{ and } S^* \not\vdash \perp; \\ -1 & \text{if } S^* \vdash \neg S'. \end{cases} \quad (3.3)$$

Fitelson's definition of k differs from the Kemeny–Oppenheim measure F (as defined above) in the additional two constraints. The reason for these constraints is that F is not defined for cases in which S' has probability 0 or 1.

Clearly, Fitelson's measure is an example of a measure of coherence that is a measure of mutual support and that satisfies the requirement that sets consisting of equivalent propositions have maximal coherence. In subsection 2.5.3, I showed that as a measure of coherence as mutual support it also has some very counterintuitive consequences. Nevertheless, in the example given there I presupposed that Fitelson's measure is intended to be an explication of our intuitive concept of coherence as mutual support. However, neither Fitelson (2003)

nor Fitelson (2004) purports to explicate our intuitive concept of coherence at all, let alone our notion of coherence as mutual support. Instead, his measure is simply intended to be the quantitative generalization of a logical notion and our intuitive concept of coherence may be quite irrelevant to that project. If it were truly irrelevant, then that would not threaten the main thesis of this chapter. For this chapter, like the earlier chapters, is concerned with our intuitive concept of coherence and its main thesis is that in order to be a satisfactory explication of our intuitive notion of coherence, a measure of coherence cannot satisfy both the dependence and the equivalence desiderata. Nevertheless, I do not believe that our intuitive notion of coherence can be irrelevant to Fitelson's project.

In fact, when we consider possible rationales for Fitelson's desiderata, it becomes apparent that he needs more than the concept of generalized equivalence in order for all of his desiderata to be acceptable. To see this, let us review each of the (subsets of) desiderata in turn.

For C there does not seem to be any problem: it seems at least quite natural that if $c(R, R')$ gives the coherence of R with R' , the overall coherence of a theory can be calculated by taking a (possibly weighted) average of the values of $c(R, R')$ for all disjoint subsets R and R' . Next, the equivalence desiderata follow from the view of coherence as generalized equivalence in the same way as the entailment desiderata follow from the view of confirmation as generalized entailment: again, the extremities should be given by, in this case, equivalence and inconsistency.

But why would we accept the dependence desiderata as desiderata for a measure of coherence? I noted above that for the project of measuring confirmation the dependence desiderata follow quite straightforwardly from the definition of probabilistic (dis)confirmation, but we are after a measure of coherence now and there is no similar standard probabilistic definition of (in)coherence. Thus, there is no straightforward manner in which the dependence (or any other) desiderata can be introduced. In this connection it is telling that the only other measures that satisfy the equivalence desiderata – Olsson's (2002) measure of coherence and Bovens and Hartmann's (2003a) difference function – both violate the dependence desiderata.⁴

It seems that in order to defend the dependence desiderata for his measure of generalized equivalence, Fitelson needs a conception of generalized equivalence as mutual support. But this conception of mutual support does not follow directly from the concept of generalized equivalence in the way the equivalence desiderata do. Instead, Fitelson needs our intuitive conception of *coherence* as mutual support in order to defend his view of a measure of generalized equivalence as a measure of mutual support. It thus appears that Fitelson cannot escape referring to our intuitive concept of coherence, even if he does not need to present his measure as an explication of our intuitive concept of coherence as

⁴As I showed in subsections 2.5.2 and 2.5.4.

mutual support.⁵ Initially, this conclusion seems rather harmless. Why would Fitelson call his measure a measure of coherence if it is not meant to appeal to any of our intuitive conceptions with respect to coherence? However, the next section will show that the combination of the two subsets of desiderata will necessarily lead to a very counterintuitive consequence, from which it follows that Fitelson's measure clearly fails to be a satisfactory explication of our intuitive concept of coherence. And if I am right in arguing that he needs to present his measure as a measure of our intuitive concept of coherence, then his proposal must be rejected.

3.3 Dependence Desiderata Versus Equivalence Desiderata

Above I discussed Fitelson's proposal that a measure of coherence must satisfy both the dependence and the equivalence desiderata. I also showed that his proposal needs to be concerned with the explication of our intuitive concept of coherence. If it is presented as such, then Fitelson's desiderata are in accordance with the view of Kemeny and Oppenheim that the explication of an intuitive concept should proceed by listing enough desiderata to determine a single ordinal class of appropriate measures. In chapter 2, I indicated that I do not believe that such a project is viable for measures of coherence as mutual support, precisely for the reason that the equivalence desiderata are intuitively incompatible with the dependence desiderata. The discussion of Fitelson's measure has paved the way for the presentation of this argument.

First, a close look at the two sets of desiderata reveals that there is at least some reason to worry about the combination of the two sets of desiderata. For consider two tautologies T_1 and T_2 . What should the coherence be of the set $\{T_1, T_2\}$? According to the equivalence desiderata, it should be maximal: they are equivalent. But according to the dependence desiderata, it should be neutral: $p(T_1 | T_2) = p(T_1)$ and $p(T_2 | T_1) = p(T_2)$. So the desiderata as defined above are actually inconsistent. As Fitelson has remarked (personal communication), however, the inconsistency can easily be removed. One obvious way to do this is by replacing d_2 by

$$(d'_2) \ c(R, R') = 0 \text{ if } p(R | \neg R') = p(R),$$

⁵There seems to be at least some support for this position in Fitelson's work. For example, Fitelson (2003: 196–197) argues that his measure is a better alternative than Shogenji's (1999) measure and Fitelson (2005c) argues that it is not too different from Bovens and Hartmann's proposal, both of which are presented as measures of coherence and both of which are explicitly linked to our coherence intuitions. If Fitelson's measure was not intended to be a competitor to these measures of coherence, then no such comparisons would be sensible.




Figure 3.1: Diagrams of the probability distributions corresponding to situations I (left) and II (right)

since $p(R \mid \neg R')$ is not defined if $p(R') = 1$. Although this move may seem somewhat *ad hoc*, I believe anyone would agree that the desiderata d_1 , d'_2 and d_3 can still be considered as desiderata for coherence as mutual support. Unfortunately, this solution only applies to cases concerning necessary truths, while the real problem is much wider. Indeed, in some cases concerning propositions with a very high probability, the equivalence desiderata and the dependence desiderata pull in opposing directions too, which leads to some very counterintuitive results. To see what I mean consider the following example.

Suppose there is a small island somewhere in the Pacific Ocean which has a population of 102 rabbits. Consider two different situations. In the first situation, 100 of the rabbits are grey and have two ears, one rabbit has two ears but is an albino, and one rabbit is grey but misses one ear. In the second situation, again 100 of the rabbits are grey and have two ears, but now the remaining rabbits are both one-eared albinos.

Next, consider the set of propositions $S = \{G, E\}$ with

G : This rabbit is grey;

E : This rabbit has two ears.

What difference should there be in the degree of coherence of S in the two situations (which are represented in figure 3.1)? Well, intuitively it seems that the set in the second situation is more coherent than in the first situation: the non-overlapping areas are smaller in the second situation while the overlapping area is equal in both cases. Intuitively, the satisfaction of these two facts seems to constitute a general condition under which a set is more coherent than another. But secondly, it also seems that the difference should not be very large: after all,




Figure 3.2: Diagrams of the probability distributions corresponding to situations III (left) and IV (right)

the non-overlapping areas have decreased only a little. But whereas Fitelson's measure agrees with the first intuition, it starkly disagrees with the second: on his measure the difference between the two situations is 1.005 on a scale of -1 to 1 : the set goes from being (slightly) incoherent to being maximally coherent.

Strikingly, this result does not depend on the number of rabbits that share both characteristics; it only depends on the number of rabbits that do not share one of them. To appreciate this consider a much larger island with 10 million rabbits. In situation III, again all rabbits are grey with two ears, except for two rabbits: one being an albino rabbit with two ears and the other being a grey one with one ear. Contrasting this with a fourth situation in which both of the remaining rabbits are one-eared albinos (the probabilities can be read off from figure 3.2), it again results that according to Fitelson's measure the difference between the two situations is slightly more than 1. More generally, it can be said that as long as in the first situation there is precisely one grey one-eared rabbit and one albino two-eared rabbit and in the second situation there are two albino one-eared rabbits, the number of grey rabbits with two ears may be arbitrarily large. The larger the number of grey rabbits, the more closely the difference in coherence between the two sets will approximate 1, but it will never become equal to or smaller than 1.

Clearly, this is a highly unsatisfactory consequence of Fitelson's measure of coherence. What may be initially less clear is that it is a necessary consequence of the combination of the dependence desiderata and the equivalence desiderata. To see that this is indeed the case, first note that in situations II and IV the two propositions are equivalent relative to the background knowledge: if the rabbit

is grey it has two ears, and if it has two ears it is grey. Therefore, coherence must be maximal. In situations I and III, by contrast, the set consists of subcontrary propositions. In section 2.6, I noted that subcontrary sets consisting of two propositions must always be any–any incoherent. Therefore, in the first and third situation the coherence must be smaller than zero. Thus, if C is to be some sort of averaging procedure, it follows that the difference in coherence in the above examples must always be larger than the difference between neutral and maximal coherence.

It seems safe to conclude that all measures of coherence that are based on both the dependence and the equivalence desiderata will necessarily have some very counterintuitive consequences. For Fitelson’s measure, this implies that coherence as a quantitative generalization of logical equivalence cannot be a measure of mutual support and for the project of explicating coherence in general this means that an explication of coherence must either reject the equivalence desiderata or reject the view of coherence as mutual support. In chapter 2, I pursued coherence as mutual support. In the next section I will pursue the equivalence desiderata by trying to find an alternative desideratum that can be combined with them in order to yield an alternative measure of coherence (which, preferably, could also function as a substitute for Fitelson’s generalization of logical equivalence). The measure that results from this will appear to formalize the intuition that lies at the basis of Siebel’s and Bovens and Hartmann’s critiques of Fitelson’s measure (see section 2.6), namely, the intuition that coherence should have something to do with relative overlap.

3.4 Measuring Coherence as Relative Overlap

In chapter 1, I introduced the concept of coherence as an epistemic virtue as having something to do with hanging together and as being somewhere in between consistency and equivalence. I argued there that there are at least two different interpretations of the ‘hanging together’ property: coherence as mutual support and coherence as set-theoretical overlap. Although the first conception has been the one most commonly appealed to in the literature on coherence, it appeared especially in section 2.6 that at least some authors have made implicit references to intuitions of relative overlap.

Above I argued that the equivalence desiderata cannot be combined with a view of coherence as mutual support. Evidently, this does not imply that they cannot be combined with a view of coherence as an epistemic virtue. Only if they could not be combined with a view of coherence as relative overlap would the conclusion follow. However, it is easy to see that this is not the case. Instead, one can construct an overlap desideratum from which the equivalence desiderata

follow as consequences. Since relative overlap is maximal in case of equivalence and minimal in case of inconsistency, a desideratum to the effect that $c(R, R')$ should depend on the relative overlap of R and R' would directly guarantee that equivalent (subsets of) propositions have maximal coherence and inconsistent (subsets of) propositions have minimal coherence.

How should we formalize such a desideratum? Arguably, relative overlap differs from support in being a property of a set of propositions instead of a relation between (sets of) propositions. While there is a clear difference between the support R lends to R' and the support R' lends to R , there is, intuitively, no similar difference between the relative overlap of R and R' on the one hand and R' and R on the other. Given this, the following desideratum seems to make the most sense:⁶

Overlap Desideratum

- (o) $c(S)$ increases as the relative overlap between the propositions in S increases,

with relative overlap defined as

$$h(S) =_{df} \frac{p(S)}{p(\bigvee S)}, \quad (3.4)$$

with S a set with two or more propositions. A concomitant result of considering overlap as a property of a subset is that the coherence of a set should be measured by averaging over the relative overlap of all its different subsets. To formalize this, let $[S]_1$ indicate the set of all subsets of S with cardinality greater than 1. Let $[[S]]_1$ denote the cardinality of $[S]_1$.⁷ Then the *overlap measure of coherence* is defined as follows:

⁶If one disagrees with this argument and feels, to the contrary, that relative overlap, too, should be formalized as a relation between (sets of) propositions, then another definition of relative overlap is required. An obvious alternative would be

$$h'(R, R') =_{df} \frac{p(R \wedge R')}{p(\bigvee R)}.$$

I am not sure whether this approach would differ much from the one proposed below, although of course the summation will be over a much larger number of elements and would therefore be much more intricate than the approach proposed below. As such, it would fail to satisfy Carnap's requirement that an explication should be as simple as possible (discussed in section 1.2). However, in the end it may turn out that my proposal below has a number of unsatisfactory consequences that the more intricate approach would help to solve. In that case, of course, we would have a good reason to pursue the latter method.

⁷For a set $S = \{R_1, \dots, R_n\}$ it holds that $[[S]]_1 = 2^n - n - 1$ (the number of the elements of the powerset of S minus the number of sets with only one element and minus 1 for the empty set).

Definition 3.2 Given a set $S = \{R_1, \dots, R_n\}$ and an ordering $\langle \hat{S}_{1_1}, \dots, \hat{S}_{1_{\llbracket S \rrbracket_1}} \rangle$ of the members of $\llbracket S \rrbracket_1$, the measure of overlap coherence of S is given by the function

$$\mathcal{R}(S) =_{df} \frac{\sum_{i=1}^{\llbracket S \rrbracket_1} h(\hat{S}_{1_i})}{\llbracket S \rrbracket_1}, \quad (3.5)$$

with $h(S)$ the relative overlap of S .

Contrary to the measures of coherence as mutual support, it is not directly clear which value would constitute neutral coherence. Whereas for a measure of coherence as mutual support, sets consisting of independent propositions have neutral (zero) coherence, the range of this measure is $(0, 1)$ and there seems to be no clear equivalent to the zero coherence in the former case. However, of course we can still compare sets to see which one is more coherent in the sense of relative overlap.⁸

This measure is clearly an explication of coherence as relative overlap. As such it does justice to the intuitions in the rabbit example above and to both Siebel's and Bovens and Hartmann's intuitions with respect to the examples discussed in section 2.6 of the last chapter.

Firstly, in the rabbit example above (section 3.3) I argued that the differences with respect to degree of coherence between situations I and II and between situations III and IV should not be too large. This intuition is respected by measure \mathcal{R} : the difference in coherence as measured by \mathcal{R} is .02 for the example with 100 rabbits and .0000002 for the example with 10 million rabbits, both on a scale of 0 to 1.

Next, Siebel's example (see section 2.6.1) concerns a subcontrary set consisting of two propositions with a relatively large overlap: out of ten murder suspects, six had committed two earlier criminal acts, while 4 had committed only one criminal act. For this example, $\mathcal{R} = .6$, which seems to accord well with Siebel's intuition that the set must be quite coherent. Finally, this measure also gives the answer to Bovens and Hartmann's example (which I have discussed in section 2.6.2) that they believe is the intuitively correct one. According to them, if in a set $S = \{R_1, R_2\}$ we set $p(R_1 \wedge R_2) = x$ and both $p(R_1 \wedge \neg R_2) = p(\neg R_1 \wedge R_2) = .05$, then increasing x from .01 to .8 should monotonically increase the coherence of the set. As graph 3.3 shows, $\mathcal{R}(S)$ satisfies this requirement.

But measure \mathcal{R} also seems better suited for analyzing coherence as a quantification of generalized equivalence. For although from that standpoint the overlap desideratum may appear just as unfounded as the dependence desiderata in Fitelson's approach, the former agrees much better with the equivalence desiderata.

⁸Alternatively, we could simply pick a value in the interval $(0, 1)$ as being that neutral value. Also, one could pick an ordinal equivalent measure that has range $(-1, 1)$ so that 0 would again constitute neutral coherence. An ordinal equivalent to h with that range is discussed below in section 3.7.




Figure 3.3: $\mathcal{R}(S)$ for values of $.01 \leq x \leq .9$

And thus, if one would accept Fitelson's argument that equivalence and inconsistency should constitute the extremities of a measure of generalized equivalence, the overlap measure would appear to be a much better candidate for coherence as generalized equivalence than Fitelson's proposal.

3.5 Comparing \mathcal{R} With the Extant Measures of Coherence

So far in this chapter, I have proposed a measure of coherence as relative overlap. But just as in the previous chapter it might have turned out that one of the extant measures of coherence was more satisfactory as a measure of coherence as relative overlap than C_d , in the case of coherence as relative overlap too, it may be the case that \mathcal{R} is less satisfactory *qua* measure of relative overlap than any of the extant accounts. Arguably, any measure of coherence as relative overlap should accord both situation III *and* situation IV a high degree of coherence (both have a very large overlap). This directly excludes Fitelson's and Shogenji's proposals and the measures C_d , C_r and C_l introduced in chapter 2. Thus, the only two measures that appear eligible are Olsson's (2002) proposal and Bovens and Hartmann's (2003a) difference function.

Of these two, Olsson's measure seems to be the most obvious option for a measure of coherence as relative overlap, since it measures the degree to which all propositions in a set overlap (see subsection 2.5.2). However, just as Shogenji's measure is not sensitive to relations of the any-any coherence type, Olsson's measure is not sensitive to relations of relative overlap other than the ones between all propositions. As an example by Bovens and Hartmann (2003a: 44–45, 50) shows, this has some untoward consequences. Consider the sets $S = \{B, G\}$ and $S' = \{B, G, P\}$ with




Figure 3.4: Diagrams of the probability distributions for S (left) and S' (right)

- B : Our pet Tweety is a bird;
 G : Our pet Tweety is a ground dweller;
 P : Our pet Tweety is a penguin.

Assume that relative to the background knowledge P entails both B and G and assume further that the probability models for the two sets are such that B and G are individually quite probable and conditionally on each other highly improbable. More specifically, assume that the probability models for both sets are given by the diagrams in figure 3.4. As Bovens and Hartmann note, set S' is intuitively more coherent than set S . However, on Olsson's account, both sets have equal coherence: $\mathcal{O}(S) = \mathcal{O}(S') = 1/99$. This, as Bovens and Hartmann rightly remark, is a very counterintuitive result.

Unsurprisingly, according to measure \mathcal{R} , in the second situation set S is more coherent than the first. The reason for this is, of course, that in the second set some of the subsets have a higher relative overlap than the two propositions in the first set. Thus, the average overlap increases, just as we would have expected. Here one may counter that the increase in coherence is less than we would have expected: $\mathcal{R}(S) \approx .01$, while $\mathcal{R}(S') \approx .015$. But from the standpoint of coherence as relative overlap, this is precisely what should be expected: the propositions that Tweety is a penguin and that Tweety is a bird do not overlap very much and similarly for the other two subsets. Furthermore, note that the difference in this situation is not fixed by the desiderata as is the case for Fitelson's measure in the rabbit example. For example, one may also choose to take the natural logarithm of \mathcal{O} , in which case the difference between the two situations becomes much larger.⁹

⁹In that case, maximal coherence would equal 0, minimal coherence would equal $-\infty$. I will not




Figure 3.5: Diagram of the new probability distribution for S'

Somewhat ironically, slightly changing Bovens and Hartmann's own example against Olsson yields an equally devastating argument against their own account. For let us assume that some (but very few) people hold ostriches as pets. This will change the probability model, for then the probability that our pet is a non-flying bird that is not a penguin is no longer zero. Suppose that it is one in a million (see figure 3.5). Then, obviously, it should still be the case that adding proposition P to set S results in a more coherent set. However, from figure 3.6 it follows that on Bovens and Hartmann's account it is indeterminate which of the two sets is the more coherent. Therefore, it seems that Bovens and Hartmann's proposal is just as unsatisfactory as Olsson's measure with respect to the question of explicating coherence as relative overlap.

3.6 Toward a Single Measure of Coherence as an Epistemic Virtue

Until now, I have presented the concept of coherence as an epistemic virtue as being based on two different intuitions: mutual support and relative overlap. By means of an example about which we have very clear coherence intuitions, I have shown that no single measure of coherence can be a satisfactory explication

here delve into the question of the desirability of a measure with a range of $(-\infty, 0)$, especially because many other scales are possible also.





Figure 3.6: $f_r(S, S')$ for $0 < x < .0005$

of both coherence as mutual support and coherence as relative overlap. Nevertheless, the above argument does not show that no measure of coherence can be satisfactory in the sense of always giving the intuitively correct answer to the question of which of two sets has the highest coherence in the sense of coherence as an epistemic virtue. The reason for this is that although it is clear that a measure of coherence cannot be both a measure of mutual support and a measure of relative overlap, it cannot be directly excluded that there exists a measure that gives the intuitively correct judgement in all of the test cases. That is, it cannot be excluded that for all genuine test cases there is a measure that does justice to our mutual-support intuitions when these are the overriding intuitions and to our relative overlap intuitions if those have overriding force.

And here, finally, Bovens and Hartmann's measure may be vindicated after all. For although I have shown that it is an unsatisfactory explication of coherence as mutual support and of coherence as relative overlap, it may be that their measure can be considered as taking the middle ground between these. Such considerations are reinforced by their own view that coherence relations are a matter both of positive relevance relations and of relative overlap relations between the propositions in a set (Bovens and Hartmann 2003a: 53). I believe that especially the Tweety example presented in the last section would give us reason to doubt that this is indeed the case. However, there is an example which makes even more clear that their measure cannot be interpreted as a measure of coherence as an epistemic virtue.

For it is evidently the case that if coherence is both a matter of overlap and of positive relevance, and of nothing else, then it should be the case that if a set S has both a higher relative overlap and a higher degree of positive relevance than another set S' , S should have a higher degree of coherence than S' . However, on Bovens and Hartmann's account this is not the case.

Figure 3.7: Diagram of the new probability distribution for S'


For consider the following example, which is a variant of a similar example discussed by Bovens and Hartmann (2003a: 39–40). In the example, a murder has been committed in Tokyo but the corpse is still to be found. A 100-square grid has been drawn over the map of the city and initially the corpse is equally likely to be found in any of the squares. Now consider two situations in each of which we receive reports about the location of the corpse from two independent and equally but only partially reliable witnesses. In the first situation, witness 1 reports that the corpse is somewhere in squares 14 to 60 (call this report ' R_1 '), and witness 2, that it is somewhere in squares 31 to 78 (R_2). In the second situation, witness 1 reports that the corpse is somewhere in squares 33 to 50 (R'_1), and witness 2, that it is somewhere in squares 50 to 67 (R'_2).

Figure 3.7 makes it easy to verify that for S we have $a_0 = .3$, $a_1 = .36$ and $a_2 = .34$, and for S' we have $a'_0 = .01$, $a'_1 = .34$ and $a'_2 = .65$. So the difference function for the two sets is this:

$$f_r(S, S') = \frac{.3 + .7(1-r)^2}{.3 + .36(1-r) + .34(1-r)^2} - \frac{.01 + .99(1-r)^2}{.01 + .34(1-r) + .65(1-r)^2}. \quad (3.6)$$

And, as figure 3.8 shows, the graph of this function crosses the r -axis. Hence, according to Bovens and Hartmann there is no fact of the matter as to which of S and S' is more coherent.

To see that this is a true counterexample against the proposal to use Bovens and Hartmann's theory as a theory of coherence as an epistemic virtue, first note that the propositions in S (greatly) support one another— $p(R_1) = p(R_2) = .48 < p(R_1 | R_2) = p(R_2 | R_1) = .625$ —while those in S' (greatly) undermine one

Figure 3.8: Graph of $f_r(S, S')$

another— $p(R'_1) = p(R'_2) = .18 > p(R'_1 | R'_2) = p(R'_2 | R'_1) \approx .056$. And since it is equally clear that the propositions in the first set overlap much more than they do in the second set, it is both the case that set S is any-any coherent while set S' is any-any incoherent and that the propositions in S have a much higher relative overlap than those in S' . Therefore, set S should be more coherent than set S' and since the differences for both the mutual support and the relative overlap are so large, this should certainly not constitute an indeterminate case. But, as we have seen above, it does for Bovens and Hartmann's theory.

The above result can even be generalized. For consider the following theorem (for a proof see Appendix 3 A):

Theorem 3.1 *Given two sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R'_1, \dots, R'_n\}$ such that $m, n > 1$ and a probability distribution $p(\cdot)$ for the two sets such that (1) $a_j = a'_k = 0$ for all $1 \leq j \leq m - 2$ and $1 \leq k \leq n - 2$ and (2) $a_{m-1} < a'_{n-1}$. In that case set S' cannot be more coherent than set S on Bovens and Hartmann's account.*

The reason this result is so counterintuitive is that it is independent of the values of a_0 and a'_0 . That is, a_0 may be arbitrarily small in comparison with a'_0 without S becoming more coherent than S' .

To illustrate the theorem, assume for the corpse in Tokyo example that an x -square grid (with x a large integer) has been drawn over the map of the city and, furthermore, that initially the corpse is equally likely to be found in any of the squares. In that case the probability that the corpse is in square i is $1/x$. Furthermore, assume that the following values for the a_i and a'_i hold:

S	S'
$a_0 = .2$	$a'_0 = .0000001$
$a_j = 0, \text{ for all } j \in \{1, \dots, m-2\}$	$a'_k = 0, \text{ for all } k \in \{1, \dots, n-2\}$
$a_{m-1} = .14$	$a'_{n-1} = .13$

It follows from the proof of theorem 3.1 that for all values of $m, n > 1$, set S is no more coherent than set S' . Also, decreasing the value of a'_0 even further to .00000001 or even smaller does not alter this result. Thus, we are left with a large class of counterintuitive examples against Bovens and Hartmann's theory of coherence.

In their (2005), Bovens and Hartmann have tried to defuse this type of counterexample. Put briefly, they argue that there is a third constituent of our intuitive judgements of coherence, namely, the specificity of the information. According to them, if coherence depends on positive relevance, relative overlap *and* specificity and if, furthermore, a set must outscore another set on all three points in order to be more coherent than that set, the examples above are no longer counterexamples against their account. For according to them it is the case in both corpse in Tokyo examples presented above that the information is more specific in the intuitively less coherent set, so that it performs better on that score. And while the much higher mutual support and relative overlap of the one set may intuitively make up for its slightly lower degree of specificity, this does not mean that a coherence ordering must follow suit:

Where some considerations are pulling so heavily in one direction that this compels our intuitive judgment, the slightest counterforce from a conflicting consideration can be sufficient for the criterion not to impose an ordering. (Bovens and Hartmann 2005)

Evidently, this rebuttal does not diminish the worth of these examples, since in this section I am specifically concerned with a notion of coherence as an epistemic virtue, which I have described as being a matter of both mutual support and relative overlap and of nothing else. However, their remark does suggest a third interpretation of coherence as an epistemic virtue, namely coherence as specificity. Unfortunately, they do not give a formal definition of this 'specificity' and it seems that we can only decide whether this notion would be a valuable addition to the other two explications of coherence as a theoretical virtue after Bovens and Hartmann provide us with such a definition.

Let me note in passing that it remains uncertain whether an appeal to specificity will cope with all possible counterexamples against Bovens and Hartmann's theory. For instance, the most obvious option, namely, the value of a_n (the higher it is, the more specific is the information in that set), does not work. For as figures 2.6 and 2.7 in the previous chapter have shown, there exist pairs of sets such

that for one set the positive relevance is higher, the relative overlap is higher and the value of a_n is higher than those of the other set, while the two still constitute an indeterminate case.

Unfortunately, all of the other options that I can think of are too strong. For example, one may stipulate that a set S is more specific than a set S' iff $\sum_{i=1}^{n-1} a_i < \sum_{i=1}^{n-1} a'_i$ for all $1 \leq i \leq n-1$ (assuming both sets have the same number of propositions). Arguably, it seems that if we accept Bovens and Hartmann's suggestion, it must also be the case that if a set is less specific than another set, it cannot be more coherent. But this is not the case for this proposal. For consider any probability model for two sets, each consisting of three propositions, such that for set S we have $a_0 = .1$, $a_1 = .06$ and $a_2 = .12$, and for S' we have $a'_0 = .1$, $a'_1 = .03$ and $a'_2 = .18$. It can easily be checked that according to Bovens and Hartmann's difference function, it is the case that set S' is more coherent than set S . However, according to the proposal for a definition of specificity at hand, it is the case that set S is more specific than set S' .

Finally, one could attempt to stipulate that a set S is more specific than a set S' iff $a_i < a'_i$ for all $1 \leq i \leq n-1$. However, this would be a very strong condition indeed, and there seems to be little justification for it. A further complication of this suggestion is that it allows for indeterminate cases with respect to specificity and it is not clear what a measure of coherence should do in such cases.

3.7 A Measure of Epistemic Coherence \mathcal{E}

In this section I will explore the option of measuring coherence as an epistemic virtue with the help of the measures of mutual support and relative overlap defined in this and the previous chapter. This section is highly tentative: as will appear below, there are a very large number of options to be explored, and I will, with only a few quite insufficient arguments, limit myself to considering only three of these. I invite everyone to experiment for him- or herself in order to find the most satisfactory measure.

The most obvious way to do justice to both intuitions of overlap and intuitions of mutual support would be by presenting the measure of coherence as an epistemic virtue as a (possibly weighted) average of coherence as mutual support and coherence as relative overlap. Given a measure of coherence as mutual support C_m and a measure of coherence as relative overlap R_n , the measure of total coherence $\mathcal{E}_{m,n}$ would then be given by:

$$\mathcal{E}_{m,n} =_{df} a C_m + b R_n, \quad (3.7)$$

with a and b averaging constants. The values of a and b can be adjusted according to what our intuitions say about certain examples, but I will only consider $a = b = 1/2$.

What measures should we choose for C_m and \mathcal{R}_n ? In the last chapter I proposed measure C_d as the most satisfactory measure of coherence as mutual support and in this chapter I have proposed measure \mathcal{R} as the most satisfactory measure of coherence as relative overlap. So should we not simply opt for the average of these two measures? I believe things are a bit more complicated than that. Recall that my case for C_d in section 2.4 was based on considerations of relative overlap. I argued there that this need not be a problem for a measure of mutual support: if the measures of confirmation disagree with respect to the question of whether the propositions in one set support each other more than in another, or vice versa, other considerations may be used to decide which of these measures is the most satisfactory. Nonetheless, in case we are looking for the best measure of coherence as mutual support to combine with a measure of relative overlap in order to yield a complete measure of coherence as an epistemic virtue, it may turn out that other intuitions should influence our choice of a measure of coherence.

Moreover, in order to compare the two measures on an equal basis it is preferable if they have the same range. If, for example, one of the measures has a range $[0, \infty)$ (as do C_l and C_r) while the other has a $[0, 1]$ range (as does \mathcal{R}), then a weighing procedure would make little sense. Indeed, in many cases it seems to make little sense to use an averaging procedure if *both* measures have range $[0, \infty)$. This poses an important dilemma for the project of measuring coherence as an epistemic virtue with measure $\mathcal{E}_{m,n}$. For although both \mathcal{R} and C_d have finite range, they do not have the same range ($[0, 1]$ and $(-\frac{1}{2}, 1)$, respectively). Moreover, it is not clear whether the other two measures, C_l and C_r , have any ordinal equivalents with finite range. At any rate, of the four measures discussed above, none have the same finite range.

It thus seems we have good reason to look at other measures of confirmation and relative overlap. Nevertheless, we need not diverge much from the four measures discussed above. For each of the measures of confirmation and of relative overlap discussed so far has at least one ordinal equivalent with a $(-1, 1)$ range. One of these we have already encountered: the Kemeny–Oppenheim measure is ordinally equivalent to the likelihood measure and the resulting measure C_k has the required range. Furthermore, an ordinal equivalent to the difference measure is

$$q(\langle S, S' \rangle) =_{df} 2 - 2^{1-d(\langle S, S' \rangle)} \quad (3.8)$$

$$= 2 - 2^{1-p(S|S') + p(S)}, \quad (3.9)$$

and it can be shown that the resulting measure C_q has range $(-1, 1)$. Note carefully, though, that I do not claim that C_q and C_d are ordinally equivalent. I only claim that q and d are ordinally equivalent and that C_q has the required range. The same will hold for the other measures proposed below.

Third,

$$s(\langle S, S' \rangle) =_{df} \frac{r(\langle S, S' \rangle) - 1}{r(\langle S, S' \rangle) + 1} \quad (3.10)$$

$$= \frac{p(S|S') - p(S)}{p(S|S') + p(S)} \quad (3.11)$$

is ordinally equivalent to $r(\langle S, S' \rangle)$ and the resulting measure C_s has a $(-1, 1)$ range. Finally,

$$o(S) =_{df} 2 h(S) - 1 \quad (3.12)$$

$$= 2 \left(\frac{p(S')}{p(\bigvee S')} \right) - 1. \quad (3.13)$$

is ordinally equivalent to h and

$$\mathcal{R}_o =_{df} \frac{\sum_{i=1}^{\llbracket S \rrbracket_1} o(\hat{S}_{1_i})}{\llbracket S \rrbracket_1} \quad (3.14)$$

has range $(-1, 1)$.

Similarly to the definition of a family of three measures as mutual support (definition 2.5), we can now define a family of three measures of coherence as an epistemic value:

Definition 3.3 *Given a set $S = \{R_1, \dots, R_n\}$ and an ordering $\langle \hat{S}_1, \dots, \hat{S}_{\llbracket S \rrbracket} \rangle$ of the members of $[S]$ and an ordering $\langle \hat{S}_{1_1}, \dots, \hat{S}_{1_{\llbracket S \rrbracket_1}} \rangle$ of the members of $[S]_1$, the degree of epistemic coherence of S is given by the function*

$$\mathcal{E}_{m,o} =_{df} \frac{1}{2} \left(\frac{\sum_{i=1}^{\llbracket S \rrbracket} m(\hat{S}_i)}{\llbracket S \rrbracket} + \frac{\sum_{i=1}^{\llbracket S \rrbracket_1} o(\hat{S}_{1_i})}{\llbracket S \rrbracket_1} \right), \quad (3.15)$$

for $m \in \{q, s, k\}$.

Again, we are left with a family of three measures and, again, it seems we will need a 'test case' to decide which of these is the most appropriate.

I believe that the following variation of the Samurai example constitutes such a test case for $\mathcal{E}_{m,o}$. Consider again the set $S = \{J, O\}$ with

J : The murderer is Japanese;

O : The murderer owns a Samurai sword.

But now consider two new situations. In the first situation, situation α , it is assumed that the murderer may be anyone living in a city of 10,000,000 inhabitants, 7 of them being Japanese and 7 of them owning Samurai swords, and 1 of




Figure 3.9: Diagrams of the probability distributions corresponding to situations α (left) and β (right)

them both being Japanese and owning a Samurai sword. In the second situation, situation β , we consider a nearby village with 100 inhabitants, 30 of them being Japanese, 30 owning a Samurai sword, and 9 both being Japanese and owning a Samurai sword (see figure 3.9).

I would propose that intuitively the set is more coherent in the first situation than in the second. Clearly, the set is independent in situation β and any-any coherent in situation α . But more importantly, intuitively the propositions seem to hang together relatively well in the first situation. For although the overlap in both cases is not very large, in situation α this seems compensated for by the fact that on the supposition of the other proposition the probability of each proposition is raised from $.0000007$ to approximately $.14$. Therefore, I would tentatively suggest that situation α is more coherent than situation β .

The outcomes for the measures of epistemic coherence $\mathcal{E}_{m,o}$ are given in the diagram below (all values are rounded to three decimal places):¹⁰

	α	β
$\mathcal{E}_{q,o}$	-.329	-.324
$\mathcal{E}_{s,o}$.0769	-.324
$\mathcal{E}_{k,o}$.0769	-.324

Clearly, only according to $\mathcal{E}_{s,o}$ and $\mathcal{E}_{k,o}$ is it the case that situation α is more coherent than situation β and we therefore have a tentative case against $\mathcal{E}_{q,o}$. Moreover,

¹⁰The reason why they all give the same value for situation β is that in that case S is independent and therefore has zero mutual support.

the rabbit example discussed above can still function as a counterexample against measure $\mathcal{E}_{k,o}$; for although the difference between the two situations is no longer necessarily larger than 1, it is still necessarily larger than $1/2$, which remains much larger than we would intuitively feel appropriate. Therefore, measure $\mathcal{E}_{s,o}$ seems the most satisfactory measure of coherence as an epistemic virtue.

As a concluding remark, let me again emphasize the tentative character of this last section. I have limited myself to considering only one ordinal equivalent of only three measures of confirmation, and have not pursued the question of different weighing factors: either for the two measures that make up $\mathcal{E}_{m,o}$ or for the multitude of elements that make up C_m and \mathcal{R}_o . Therefore, there is much room for experiment here and those who feel unconvinced by the version of the Samurai example presented in this section are free to substitute their own favorite measure of confirmation or to fiddle around with different weighing factors.

3.8 Conclusion

This chapter has presented two different measures of coherence. From a discussion of Fitelson's view of coherence as generalized equivalence it has appeared that only a measure of relative overlap can do justice to the desiderata that follow from such a conception of coherence. Moreover, it appeared that the resulting measure \mathcal{R} does justice to coherence intuitions that are not respected by a measure of coherence as mutual support. Those who agree with the intuitions proposed by Bovens and Hartmann and Siebel concerning their examples will find that \mathcal{R} completely accords with these intuitions.

Secondly, I have presented a measure of coherence as an epistemic virtue that purports to do justice to our intuitions of coherence both as mutual support and as relative overlap. Those who have tried the resulting measure for some of the test cases of coherence (both the ones discussed in this thesis so far and others present in the literature) will have found that $\mathcal{E}_{m,o}$ gives the intuitively correct answer in many cases.

Although it would be preferable to have a single measure of coherence as an epistemic virtue, one should be aware of the limitations of such an approach. For even if $\mathcal{E}_{m,o}$ gives the intuitively correct answer in many cases, it can never be a complete measure of mutual support or of relative overlap. For example, according to $\mathcal{E}_{m,o}$ it will not be the case that an independent set is always less coherent than an any-any coherent set. Nor is it the case that $\mathcal{E}_{m,o}$ is a monotonically increasing function of x in the example that Bovens and Hartmann use to criticize Fitelson's measure of coherence (see subsection 2.6.2). Thus, in order to do full justice to our intuitions of mutual support and of relative overlap we will still require two different measures.

What has become clear, however, is that none of the extant measures of coherence is satisfactory for measuring coherence as an epistemic virtue, even if we agree that we would require two different measures to measure the two different aspects of that notion. Altogether, this may not be surprising, especially because none of the extant measures of coherence has been presented as an explication of coherence as an epistemic virtue. And while Fitelson's approach seems to encounter fatal difficulties even as a measure of generalized equivalence, I have not yet considered the other measures as measures of coherence as a confidence boosting property. To this subject I will turn in the next chapter.

Appendix 3 A: Proof of Theorem 3.1

Theorem 3.1 *Given two sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R'_1, \dots, R'_n\}$ such that $m, n > 1$ and a probability distribution $p(\cdot)$ for the two sets such that (1) $a_j = a'_k = 0$ for all $1 \leq j \leq m-2$ and $1 \leq k \leq n-2$ and (2) $a_{m-1} < a'_{n-1}$. In that case set S' cannot be more coherent than set S on Bovens and Hartmann's account.*

For it to be the case that S' is not more coherent than S , there must be at least one value r_0 such that $f_{r_0}(S', S) < 0$. Furthermore, note that for this to be the case it would suffice if the derivative of $f_r(S', S)$ is negative in $r = 0$, since this implies that $f_r(S', S) < 0$ for values of r approximating 0.¹¹

To find the expression for the derivative of $f_r(S', S)$ for $r = 0$, first consider $c_r(S^*)$ for a set $S^* = \{R_1, \dots, R_K\}$ with K propositions:

$$c(S^*) = \frac{a_0 + (1 - a_0)(1 - r)^K}{\sum_{i=0}^K a_i (1 - r)^i},$$

where I have dropped the subscript r . Next, substitute $1 - r = x$. Then:

$$\frac{dc(S^*)}{dr} = -\frac{dc(x)}{dx}.$$

For r approximating 0, x approximates 1. Next, calculate $(\frac{dc(x)}{dx})_{x=1}$ using $\sum_{i=0}^K a_i = 1$:

$$\begin{aligned} \left(\frac{dc(x)}{dx}\right)_{x=1} &= \left(\left(\frac{1}{\sum_{i=0}^K a_i x^i}\right) K(1 - a_0)x^{K-1}\right)_{x=1} - \\ &\quad \left(\left(a_0 + (1 - a_0)x^K\right) \left(\frac{\sum_{i=0}^K a_i i x^{i-1}}{\left(\sum_{i=0}^K a_i x^i\right)^2}\right)\right)_{x=1} \\ &= \left(\frac{1}{\sum_{i=0}^K a_i}\right) K(1 - a_0) - (a_0 + (1 - a_0)) \left(\frac{\sum_{i=0}^K a_i i}{\left(\sum_{i=0}^K a_i\right)^2}\right) \\ &= K(1 - a_0) - \left(\sum_{i=0}^K a_i i\right) \\ &= K(1 - a_0) - \left(\sum_{i=0}^{K-1} i a_i + K a_K\right) \\ &= K(1 - a_0) - \left(\sum_{i=0}^{K-1} i a_i + K \left(1 - \sum_{i=0}^{K-1} a_i\right)\right) \end{aligned}$$

¹¹I thank Stephan Hartmann for pointing out that this approach is much easier than the one I originally had in mind.

$$\begin{aligned}
&= (K - Ka_0) - \left(K + \sum_{i=0}^{K-1} (i - K)a_i \right) \\
&= -Ka_0 - \left(\sum_{i=0}^{K-1} (i - K)a_i \right) \\
&= - \sum_{i=1}^{K-1} (i - K)a_i.
\end{aligned}$$

Since $\frac{dc(S^*)}{dr} = -\frac{dc(x)}{dx}$, it follows that

$$\left(\frac{dc(S^*)}{dr} \right)_{r=0} = \sum_{i=1}^{K-1} (i - K)a_i. \quad (3.16)$$

Now consider two sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R'_1, \dots, R'_n\}$ such that $m, n > 1$ and a probability distribution $p(\cdot)$ for the two sets such that (1) $a_j = a'_k = 0$ for all $1 \leq j \leq m-2$ and $1 \leq k \leq n-2$ and (2) $a_{m-1} < a'_{n-1}$. Using equation (3.16), it follows that the derivative of $f(S', S)$ for $r = 0$ is negative:

$$\begin{aligned}
\left(\frac{df(S', S)}{dr} \right)_{r=0} &= \sum_{i=1}^{n-1} (i - n)a'_i - \sum_{i=1}^{m-1} (i - m)a_i \\
&= ((n-1) - n)a'_{n-1} - ((m-1) - m)a_{m-1} \\
&= a_{m-1} - a'_{n-1} \\
&< 0
\end{aligned}$$

Thus, set S' is no more coherent than set S . ■

Chapter 4

Coherence as a Confidence Boosting Property

4.1 Introduction

In this chapter I will explicate the notion of coherence as a confidence boosting property. In the introductory chapter I discussed how Lewis and Bonjour have introduced coherence as part of a rebuttal to the epistemic skeptic, who argues that we have no justification for any (or most) of our knowledge claims. According to Bonjour, if we are presented with a number of reports by independent, highly unreliable witnesses, we may still be justified in accepting each of these reports if they are sufficiently coherent. That is, if the information we are presented with is very coherent, then it is much more likely to be true. Evidently, this assumes that coherence is truth conducive.

This view has been challenged by, among others, Klein and Warfield (1994) and (1996). According to them, adding a proposition to a set of propositions can intuitively increase its coherence while it can never increase its probability (at best, the probability will remain the same). Their argument has led to the first tentative proposals of measures of coherence by Shogenji (1999) and Olsson (2002). Although all of these authors have taken for granted the model of witnesses reporting on propositions (or sentences), none of them have taken these reports themselves into consideration. Instead, they regard only the probability of the propositions in a set and do not take into consideration the probability we would assign to these propositions if witnesses with a given reliability had reported on them.¹

¹This remark no longer applies to Olsson's (2005) paper, more on which below (chapter 6).

In their *Bayesian Epistemology*, Bovens and Hartmann improve on these discussions by explicating the notion of coherence as a confidence boosting property, where the confidence boost is the increase in our confidence that the information is true due to witness reports to that effect. In taking this view, they present the concept of *Bayesian Coherentism* as a formal description of the view of coherence as being truth conducive. From their discussion of Bayesian Coherentism it follows that no theory of coherence could ever aspire to yield a complete coherence ordering on all sets of propositions. The best we can hope for is a *quasi-ordering*, that is, an ordering which is reflexive and transitive, but not necessarily complete. I will discuss this claim in section 4.2, while in section 4.3 I will present their theory of coherence as a confidence boosting property.

Although their theory of coherence clearly is a great improvement on its predecessors, sections 4.4 and 4.5 will show that it runs into some serious trouble. Fortunately, this trouble may be relieved if we remove a small element from their theory of coherence, namely (what I will call) the *maximality requirement*. This requirement is highly reminiscent of Fitelson's equivalence desiderata (see section 3.3) and in section 4.6 I will show how abandoning it will enable us to formulate a new measure of coherence that is equally based on the idea of coherence as a confidence boosting property, but that does not yield the counterintuitive consequences that Bovens and Hartmann's account yields. Indeed, this revised theory of coherence as a confidence boosting property will solve almost all of the counterexamples against Bovens and Hartmann's theory presented in this and the previous chapters.

The conclusion of this chapter, therefore, will be that the examples proposed in this and earlier chapters do not necessarily challenge Bovens and Hartmann's basic intuitions.

4.2 Bovens and Hartmann's Impossibility Result

As noted above, Bovens and Hartmann follow Lewis's and Bonjour's general thrust by considering only propositions that have been reported by independent, partially reliable witnesses. To make this formally precise, they define the notion of an *information set* S as a finite set of propositions about each of which we have been informed by a separate source. These information sources are called *witnesses* and it is assumed that they are independent of each other and that they are neither fully reliable nor fully unreliable. Also, it is supposed that they supply only positive or negative reports, that is, the witnesses report that something is or is not the case; they do not report probabilities. Next, S is defined as the set of all such information sets.

As will be clear from the introduction, the question at stake in this chapter

is whether coherence can be a confidence boosting property. Or put differently, does the fact that propositions in an information set cohere increase our confidence that the information is true? To answer this question we must first determine what other factors influence our confidence that the information is true (otherwise we could never say whether the increase in confidence is due to an increase in coherence or an increase in any of the other factors).

Thus the following question arises: assuming positive reports on all of the propositions in an information set, what are the factors that determine our confidence that the conjunction of the propositions in that set is true? According to Bovens and Hartmann, there are three such factors: the prior probability that the information is true; the reliability of the witnesses; and the extent to which the propositions cohere. These factors should be expressed as *ceteris paribus* claims. The *ceteris paribus* condition for the prior probability stipulates that 'the more expected (or equivalently, the less surprising) the information is, the greater our degree of confidence, *ceteris paribus*' (*op.cit.*: 11). Similarly, an increase in witness reliability will make us more confident that the information is true, *ceteris paribus*. Finally, the *ceteris paribus* condition for coherence reads: '[w]hen we gather information from independent and partially and equally reliable sources, the more coherent the story is, the more confident we are that the story is true, *ceteris paribus*' (*op.cit.*: 31). Clearly, the latter statement expresses the view that coherence is a confidence boosting property. However, it is still rather vague. To make it more precise, Bovens and Hartmann define the notion of *Bayesian Coherentism* (henceforth BC), as the combination of BC_1 and BC_2 , which are defined as follows:

- BC_1 'For all information sets $S, S' \in \mathbf{S}$, if S is no less coherent than S' , then our degree of confidence that the content of S (i.e. the conjunction of the propositions in S) is true is no less than our degree of confidence that the content of S' is true, *ceteris paribus*.' (*op.cit.*: 11)
- BC_2 The relation of 'being no less coherent than' is an ordering and is fully determined by the probabilistic features of the information sets contained in \mathbf{S} (cf. Bovens and Hartmann 2003a: 25).

To evaluate these claims, it is helpful to derive an expression for our confidence that the content of an information set is true. That is, we need to determine the probability that the conjunction of propositions in a set S is true, given positive reports on all of these propositions.

To this effect, consider an information set $S = \{R_1, \dots, R_n\}$ and a probability distribution $p(\cdot)$ over the elements of S , and let ' a_i ' stand for the sum of the probabilities of the conjunctions consisting of $n - i$ elements of S and the negations of all the remaining elements of S . So, for instance, if $S = \{R_1, R_2\}$, then

$a_0 = p(R_1 \wedge R_2)$, $a_1 = p(R_1 \wedge \neg R_2) + p(\neg R_1 \wedge R_2)$, and $a_2 = p(\neg R_1 \wedge \neg R_2)$. Furthermore, let REPR_j be a report by witness j to the effect that R_j is the case (there is one witness for each proposition). To model the reliability of the witnesses, let

$$p_j =_{df} p(\text{REPR}_j | R_j) \quad (4.1)$$

be the true-positive rate of witness j with respect to proposition R_j and let

$$q_j =_{df} p(\text{REPR}_j | \neg R_j) \quad (4.2)$$

be the rate of false positives. Assume that all witnesses are equally reliable: for all j , we have $p_j = p$ and $q_j = q$. Next, define

$$r =_{df} 1 - q/p \quad (4.3)$$

as the reliability of the witnesses.² They are assumed to be neither fully reliable ($r \neq 1$) nor fully unreliable ($r \neq 0$); thus, $r \in (0, 1)$. With these instruments, it becomes possible to calculate the posterior probability

$$p^*(R_1 \wedge \dots \wedge R_n) =_{df} p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n), \quad (4.4)$$

which gives the probability of the conjunction of the elements of set S given positive reports with respect to all of these elements. It can be shown (see Bovens and Hartmann 2003a, Appendix A1) that if $a_0 > 0$, the following relation holds:

$$p^*(R_1 \wedge \dots \wedge R_n) = \frac{a_0}{\sum_{i=0}^n a_i (1-r)^i}. \quad (4.5)$$

According to Bovens and Hartmann, the reliability of the witnesses should not influence the coherence of an information set. This seems to make sense, since we would not want to adopt a model in which an information set becomes more coherent if the witnesses become more reliable. Given this assumption, BC will be refuted if we find probability distributions for two sets S and S' such that they have equal prior probability and such that there are values of the reliability parameter r_1 and r_2 such that our confidence in set S is higher than our confidence in set S' given reliability r_1 and our confidence in set S' is higher than our confidence in set S given reliability r_2 . For in both cases ($r = r_1$ and $r = r_2$) there is a difference in the posterior probability p^* of both sets that is not caused by differences in prior probability or witness reliability (since both are fixed in both cases). Since the posterior probability is different for the two cases, it follows that there are only three options left (if we want to save the concept of coherence as confidence boosting): (1) coherence depends on the reliability of the witnesses;

²Bovens and Hartmann (*op.cit.*: 22) show that the precise way in which the reliability of the witnesses is modeled does not affect their impossibility result.




Figure 4.1: Graph of $p^*(S) - p^*(S')$

(2) there is a fourth factor that determines our confidence in the information set; or (3) the coherence ordering is not complete. If we find (1) counterintuitive and if we agree with Bovens and Hartmann (*op. cit.*: 21–22) with respect to (2) that it is completely unclear what this fourth factor might be, then we are left with the conclusion that BC_2 is false.

But it is easy to see that such cases exist. For example, all probability distributions for sets $S = \{R_1, R_2, R_3\}$ and $S' = \{R'_1, R'_2, R'_3\}$ for which

S			S'		
a_0	=	.05	a'_0	=	.05
a_1	=	.30	a'_1	=	.20
a_2	=	.10	a'_2	=	.70

have the required characteristics for a counterexample (see Bovens and Hartmann 2003a: 20). The difference between the posterior probabilities of the sets for all values of $0 < r < 1$ has been plotted in figure 4.1. Clearly, $p^*(S) > p^*(S')$ for $r < .8$ and $p^*(S') > p^*(S)$ for $r > .8$ and since the prior probability of both sets is equal, we have a genuine counterexample against BC. Thus if we accept that coherence is not influenced by the reliability of the witnesses and if we believe that there is no fourth factor that determines p^* , it follows that BC is refuted. What does this mean for Bayesian Coherentism? Does it mean that the epistemic skeptic was right all along and that no matter how coherent the information we receive is, this does not give us any reason to be confident that it is true? Well, contrary to the claims of Klein and Warfield, this proof does not purport to show that no form of Bayesian Coherentism is possible. Instead, Bovens and Hartmann (2003a: §1.5) show that if BC_2 is replaced by

BC'_2 The relation of ‘being no less coherent than’ is a quasi-ordering and is fully

determined by the probabilistic features of the information sets contained in S ,

the impossibility result no longer goes through (recall that a quasi-ordering is an ordering that is reflexive, transitive, but not necessarily complete). In that case the sets S and S' used in the derivation of the impossibility result above could simply be an example of an indeterminate case, to wit, a case in which there is no coherence ordering possible. And, as will appear below, this is precisely what Bovens and Hartmann's theory of coherence will tell us about these two sets. For future reference, define BC^* as the conjunction of BC_1 and BC'_2 .

4.3 Bovens and Hartmann's Theory of Coherence

Bovens and Hartmann define the *probability boost* $b(R_1 \wedge \dots \wedge R_n)$ as the ratio of the posterior probability and the prior probability:

$$b(R_1 \wedge \dots \wedge R_n) = \frac{p^*(R_1 \wedge \dots \wedge R_n)}{p(R_1 \wedge \dots \wedge R_n)}. \quad (4.6)$$

If coherence is to be a confidence boosting property, then it must be the case that if the coherence increases, the value of b increases also. As noted above, Bovens and Hartmann believe that there are three different factors that influence the probability boost: witness reliability, prior probability, and coherence. To find a true measure of coherence, therefore, one needs to distinguish the degree to which coherence contributes to the value of $b(R_1 \wedge \dots \wedge R_n)$ from the contributions of witness reliability and prior probability. But it is not straightforward how this may be achieved. The basic problem is that although reliability may be convincingly portrayed as being independent from coherence, this does not hold for the prior probability. To see this, remember that the prior probability of the information set is the prior probability that all the elements of the information set are true, or, in other words, the measure in which they overlap. In each of the examples presented in earlier chapters, the extent to which the propositions in a set overlap seemed to play an important role in our intuitive judgements of coherence.

Therefore, neither BC nor BC^* will be of much help in determining a measure of coherence. If the element that plays an important role for coherence (to wit, the prior probability) must be kept constant for the definition to apply, then we will have no guidance on how to measure the contribution of the prior probability to the coherence. An additional element or criterion is required to disentangle the prior probability from the coherence of a set.

The element that Bovens and Hartmann add and which I will call the *maximality requirement* is highly reminiscent of the equivalence desiderata discussed in

chapter 3. According to them a set is maximally coherent if it consists of equivalent propositions (*op. cit.*: 32–35), where two propositions R_i and R_j are equivalent if they are logically equivalent relative to the background knowledge. Call such a set an *equivalent set* and call a set *nonequivalent* if not all of its propositions are logically equivalent relative to the background knowledge. Furthermore, let two sets of propositions S_i and S_j be equivalent if the conjunctions of the propositions in the sets are equivalent relative to the background knowledge. According to Bovens and Hartmann, the property of maximal coherence should not depend on a set's prior probability or the number of propositions. Using this requirement, they define the notion of a maximal confidence boost, i.e., the confidence boost a set would have received if the information set had consisted of equivalent propositions, in other words, if all of a_1, \dots, a_{n-1} had been zero (Bovens and Hartmann 2003a: 32–33). In that case $a_n = 1 - a_0$,³ and so equation (4.5) reduces to

$$p_{\max}^*(R_1 \wedge \dots \wedge R_n) = \frac{a_0}{a_0 + (1 - a_0)(1 - r)^n}. \quad (4.7)$$

Next, one can define a measure $c_r(S)$ as the fraction of the actual boost over the maximal boost (for a derivation, see Bovens and Hartmann 2003a: 34):

$$c_r(S) =_{df} \frac{b(R_1 \wedge \dots \wedge R_n)}{b_{\max}(R_1 \wedge \dots \wedge R_n)} = \frac{a_0 + (1 - a_0)(1 - r)^n}{\sum_{i=0}^n a_i(1 - r)^i}. \quad (4.8)$$

It is easy to see that $c_r(S) = 1$ if R_1, \dots, R_n are equivalent, and $c_r(S) < 1$ if otherwise. Next, the *difference function* calculates the difference between the values of c_r for two different sets:

$$f_r(S, S') = c_r(S) - c_r(S'). \quad (4.9)$$

Of course, this function still depends critically on the reliability of the witnesses. All else equal, a higher value for r will raise the value of $f_r(S, S')$. As announced above, Bovens and Hartmann believe that coherence should be independent of witness reliability, and they cancel out the dependence on the reliability by defining the two-place relation of *being no less coherent than*, ' \succcurlyeq ', as follows:

$$\text{for all } S, S' \in \mathbf{S}, \quad S \succcurlyeq S' \quad \text{iff} \quad f_r(S, S') \geq 0 \quad \text{for all } r \in (0, 1). \quad (4.10)$$

Interestingly, on this account there is no such thing as the coherence of a set: it is impossible to say whether a set is (in)coherent. Instead, only comparative judgments are possible: a set S is no less coherent than another set S' iff $f_r(S, S') \geq 0$ for all values of the reliability parameter.⁴ Also, since $f_r(S, S')$ may be positive

³Since $\sum_{i=0}^n a_i = 1$ and in this case $\sum_{i=1}^{n-1} a_i = 0$, we have $a_0 + a_n = 1$.

⁴One may feel that the introduction of the ' \succcurlyeq '-relation as 'being no less coherent than' is rather awkward also. Would we not want our theory of coherence to produce a ' $>$ ' relation of 'being more coherent than'? But this can easily be remedied by stipulating that $S > S'$ iff $f_r(S, S') > 0$ for all values of $r \in (0, 1)$ and $f_r(S, S') > 0$ for at least one value of $r \in (0, 1)$.

for some values of r and negative for others, this analysis yields only a quasi-ordering: there may be sets S and S' such that neither $S \succcurlyeq S'$ nor $S' \succcurlyeq S$. This is in accordance with the impossibility result. Indeed, it can be checked that the sets S and S' that were used in the impossibility result do in fact constitute such an indeterminate case.

More generally, it can easily be seen that in all cases where there is an impossibility result like the one above Bovens and Hartmann's difference function will not yield an ordering. For consider two sets S and S' with an equal number of propositions and equal prior probability. In that case, the maximal confidence boost will be equal for both sets and $f_r(S, S')$ becomes linearly dependent on the difference between $p^*(S)$ and $p^*(S')$. Therefore, for all probability models such that there are values of the reliability parameter r_1 and r_2 such that our confidence in set S is higher than our confidence in set S' given reliability r_1 and our confidence in set S' is higher than our confidence in set S given reliability r_2 , it will be the case that $f_r(S, S')$ is positive for values of $r = r_1$ and negative for $r = r_2$ and thus no coherence ordering is possible. However, as the next two sections will show, a different problem emerges if we drop the condition that the number of propositions remains equal.

4.4 A Different Kind of Impossibility Result

In chapters 2 and 4 above, I presented various examples in which Bovens and Hartmann's measure behaves counterintuitively. For instance, in subsection 2.5.4 it appeared that independent sets are not generally less coherent than positively dependent sets, even if only in a *ceteris paribus* sense. Furthermore, in section 3.6 of the last chapter it appeared that in cases where the propositions in one set have both a substantially larger relative overlap than in another set and support each other to a substantially higher degree, the first set is not necessarily more coherent than the second. It seems quite natural to claim that in such cases it does not really matter what notion of coherence we are trying to explicate. Instead, such examples seem to be devastating to all accounts that attempt to explicate our intuitive concept of coherence. Nonetheless, given the above impossibility result, Bovens and Hartmann may seem to have an easy reply to such examples. For if it is true that coherence cannot unqualifiedly increase our confidence in an information set, then it may seem not so awkward that in some specific examples it is indeed the case that it is indeterminate whether one set is more coherent than another.

Unfortunately, this answer cannot work in all cases. That is, I believe that an argument can be made to the effect that in some cases one set *should* be more coherent than another according to Bovens and Hartmann's discussion of what

it means for a property to be confidence boosting. And as will appear in this and the next section, the difference function discussed above does not give the correct verdict in all of these cases.

To see how such an argument may proceed, recall that Bovens and Hartmann's impossibility result shows us that there can be no complete coherence ordering. In the example two sets with equal prior probabilities are compared, but it turns out that for some values of the reliability parameter one set is more probable given positive reports by all of the witnesses, while for other values of the parameter the other set is more probable. And since the two other determinants of our posterior probability are equal in both cases, the differences between the posterior probabilities have to be due to differences in coherence.

Clearly, this result tells us little if anything about cases in which the prior probability is not fixed in a similar fashion. For instance, a theory of coherence may yield an indeterminate verdict even if the posterior probability of one set is higher than the posterior probability of another for all values of the reliability parameter. For it may simply be the case that the difference in posterior probability is caused by the fact that the prior probability of the first set is higher than that of the second. However, the impossibility result does tell us that if two sets have the same prior probability, while one of the sets has a higher posterior probability than the other for all values of the reliability parameter, then a theory of coherence must classify the first set as more coherent than the second. For in that case the difference in posterior probability cannot be due to differences in either the prior probability or the reliability of the witnesses, and, since these are the only two other determinants of our confidence according to Bovens and Hartmann, the differences in posterior probability must be caused by differences in coherence.

We have already encountered one example in which this condition is satisfied, namely, the Tweety example discussed in section 3.5. Recall that in the Tweety example the proposition that Tweety is a penguin is added to a set consisting of the propositions that Tweety is a bird and that Tweety is a ground-dweller. In Bovens and Hartmann's version of the example, the value of the prior probability that the information is true is equal in both sets: $a_0 = a'_0 = .01$. Moreover, it can be shown that the posterior probability of the second set is higher than that of the first for all values of the reliability parameter. Indeed, this is a consequence of the following, more general, theorem (for a proof, see Appendix 4 A):

Theorem 4.1 *Given two sets $S = \{R_1, \dots, R_{n-1}\}$ and $S' = \{R_1, \dots, R_n\}$, and with $0 < p(R_n) < 1$ and $p(R_1 \wedge \dots \wedge R_{n-1}) = p(R_1 \wedge \dots \wedge R_n)$, it will be the case that $p^*(R_1 \wedge \dots \wedge R_n) > p^*(R_1 \wedge \dots \wedge R_{n-1})$ for all $r \in (0, 1)$.*

Therefore, in the Tweety example it should be the case that the second and larger set is more coherent than the first and, as is shown by Bovens and Hartmann

(2003a: 44–45), their theory of coherence agrees with this conclusion.

But of course the above argument does not tell us that in some cases adding a proposition to an information set should increase its coherence. No, it tells us that in *all* cases where we add a contingent proposition to an information set that does not affect its prior probability, the coherence should increase. But, as the following theorem shows, for Bovens and Hartmann's theory of coherence this will not always be the case (for a proof, see Appendix 4 A):

Theorem 4.2 *Adding a proposition R_{m+1} to a nonequivalent set $S = \{R_1, \dots, R_m\}$ will increase S 's coherence iff R_{m+1} is equivalent with $\bigwedge S$, i.e., iff $p(R_1 \wedge \dots \wedge R_m \mid R_{m+1}) = p(R_{m+1} \mid R_1 \wedge \dots \wedge R_m) = 1$.*

Evidently, this theorem shows (as we have seen in their original Tweety example) that on Bovens and Hartmann's theory of coherence it may be the case that adding a proposition to a set may sometimes increase its coherence. But on the other hand it also shows that the coherence cannot increase *unless* the proposition we add is equivalent to the conjunction of propositions in that set. But according to theorem 4.1, as long as the prior probability does not decrease, the posterior probability always increases. Therefore, we have a large class of counterexamples against Bovens and Hartmann's theory of coherence as a confidence boosting property, to wit, all cases in which we add a contingent proposition to a set of propositions that is not equivalent to the conjunction of propositions in that set but which does not decrease its prior probability. For in all such cases the increase in posterior probability cannot be due to differences in either the prior probabilities of the sets or the reliability of the witnesses, while according to Bovens and Hartmann's difference function it cannot be due to differences in coherence either.

To illustrate this, consider the following variation of the Tweety example. In this example we compare the sets $S = \{B, F\}$ and $S' = \{B, F, A\}$ with

- B : Our pet Tweety is a bird;
- F : Our pet Tweety is of a nonflying species;
- A : Our pet Tweety's natural habitat is Antarctica.

Assume that the probability distributions for the two sets are such that B and F are individually relatively probable and conditionally on each other highly improbable. Furthermore, assume that the conjunction of B and F entails A : in this case there are no ostriches or other ground-dwelling birds whose natural habitat is not Antarctica. But now assume that there is an extremely slight probability (say 10^{-10}) that Tweety is a blue whale. More specifically, assume that the probability models for both sets are given by the diagrams in figure 4.2. Evidently, the prior probability of both sets is equal and, therefore, adding the proposition that Tweety's natural habitat is Antarctica will increase our confidence that the




Figure 4.2: Diagrams of the probability distributions for S (left) and S' (right)

information is true. Nevertheless, it can be checked that $f_r(S, S') < 0$ for values of r below, approximately, 10^{-8} .

To conclude this section, it seems that we have derived a general class of counterexamples against Bovens and Hartmann's theory of coherence as a confidence boosting property. In the next section I will discuss a number of obvious solutions to these problems.

4.5 Possible Solutions for Bovens and Hartmann

While in the previous chapters I showed that Bovens and Hartmann's theory of coherence leads to some strongly counterintuitive consequences, in this chapter I have shown that it is deficient also as an explication of coherence as a confidence boosting property.

Now the question naturally arises whether any of these results are unavoidable given the model that Bovens and Hartmann have chosen for their analysis of coherence. As became clear above, *Bayesian Epistemology* makes fundamental use of the idea of coherence as a confidence boosting property and of the concept of modeling coherence as a property of an information set with independent witnesses reporting on the propositions of that information set. If the result derived in the previous section is the necessary byproduct of such an analysis, then this approach would be highly unsatisfactory for the analysis of coherence as a confidence boosting property. And, moreover, since it is unclear whether there are other approaches to the explication of coherence as confidence boosting, it may turn out that no satisfactory explication of coherence as a confidence boosting property is possible. For all proponents of BC or BC*, this would be a

highly unsettling result.

In this section, I will argue that results like those derived in the last section are indeed the necessary consequences of Bovens and Hartmann's model. Nonetheless, I will show in the next section that this does not seriously diminish the prospects for *Bayesian Coherentism*. I will propose a new model in which coherence is still a confidence boosting property, but in which at least the most counterintuitive results no longer arise. But let me first show that the prospects for Bovens and Hartmann's own theory are not very bright.

Perhaps the most obvious way in which one may attempt to salvage *Bayesian Coherentism* is by using another measure of support to measure the probability boost. Recall from the discussion above that Bovens and Hartmann define the probability boost as the ratio between the posterior and the prior probability:

$$b(R_1 \wedge \dots \wedge R_n) = \frac{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}{p(R_1 \wedge \dots \wedge R_n)}. \quad (4.11)$$

That is, they use the ratio-measure of confirmation to measure the probability boost. But of course they could equally well have defined the probability boost by means of another measure of confirmation.⁵

Nonetheless, I do not believe that such an approach would be successful. One reason for this is that neither of the two most popular alternatives – the difference measure and the likelihood measure – would do the job.

Recall that the prior probability $p(R_1 \wedge \dots \wedge R_n)$ equals a_0 , while the posterior probability $p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)$ is given by

$$p^*(R_1 \wedge \dots \wedge R_n) = \frac{a_0}{\sum_{i=0}^n a_i (1-r)^i}. \quad (4.12)$$

Measuring the probability boost by means of the likelihood measure would lead to the following function:

$$b^l(R_1 \wedge \dots \wedge R_n) = \frac{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n)}{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | \neg(R_1 \wedge \dots \wedge R_n))}. \quad (4.13)$$

In a similar fashion to Bovens and Hartmann's derivation of their expression for $p^*(R_1 \wedge \dots \wedge R_n)$ we can rewrite this expression in terms of the reliability parameter r and the a_i . If we again take the ratio between the probability boost and the maximal probability boost, the following measure results (the derivation is given in Appendix 4 B):

$$c_r^l(S) =_{df} \frac{b^l(R_1 \wedge \dots \wedge R_n)}{b_{\max}^l(R_1 \wedge \dots \wedge R_n)} = \frac{(1-a_0)(1-r)^n}{\sum_{i=1}^n a_i (1-r)^i}. \quad (4.14)$$

⁵This strategy was suggested to me by Branden Fitelson.




Figure 4.3: Diagrams of the new probability distributions for S (left) and S' (right)

But the resulting difference function ($f_r^l(S, S') =_{df} c_r^l(S) - c_r^l(S')$) performs even less satisfactorily than the original one. For according to this function it is the case in the Tweety example discussed above that the *less* probable set is actually *more* coherent. That is, while Bovens and Hartmann's difference function classifies the case as indeterminate, this alternative difference function classifies the wrong set as more coherent.

Whereas the likelihood difference function performs less satisfactorily than Bovens and Hartmann's original measure, measuring the probability boost by means of the difference measure leads to a more satisfactory difference function. For in that case we would have the following coherence measure:

$$c_r^d(S) =_{df} \frac{b^d(R_1 \wedge \dots \wedge R_n)}{b_{\max}^d(R_1 \wedge \dots \wedge R_n)} \quad (4.15)$$

$$= \frac{\frac{a_0}{\sum_{i=0}^n a_i(1-r)^i} - a_0}{\frac{a_0}{a_0 + (1-a_0)(1-r)^n} - a_0} \quad (4.16)$$

$$= \frac{\frac{1}{\sum_{i=0}^n a_i(1-r)^i} - 1}{\frac{1}{a_0 + (1-a_0)(1-r)^n} - 1}. \quad (4.17)$$

The new difference function ($f_r^d(S, S') =_{df} c_r^d(S) - c_r^d(S')$) actually gives the correct judgements in the Tweety example. Therefore, it follows immediately that theorem 4.2 does not hold for $f_r^d(S, S')$. However, there still exist counterexamples against this measure. For consider the probability distribution for sets $S = \{B, F\}$ and $S' = \{B, F, A\}$ (defined in section 4.4) given in figure 4.3. Clearly, $a_0 = a'_0$ and (by theorem 4.1) $p^*(S) > p^*(S')$ for all values of r . Nonetheless, it can easily be checked that $f_r^d(S, S') < 0$ for values of r between, approximately,

.55 and .72 and $f_r^d(S, S') > 0$ otherwise.⁶ Therefore, the type of counterexample presented in the previous section cannot be countered by measuring the probability boost by means of either the likelihood measure or the difference measure.⁷

Evidently, the fact that the difference measure seems to perform better than the ratio measure as a measure of the probability boost may encourage one to search for another measure of confirmation to measure the probability boost. However, I feel that there are good *a priori* grounds for believing that such a project will not be fruitful unless Bovens and Hartmann's model is substantially altered. The reason for this is that I believe that Bovens and Hartmann's model involves an inconsistency.

In section 4.3, I indicated that Bovens and Hartmann need an additional element to distinguish coherence and prior probability, to wit, the *maximality requirement*. As intimated there, this requirement stipulates that equivalent sets are maximally coherent *per se*, i.e., irrespective of the value of a_0 and of the number of propositions. However, this requirement – which lies at the basis of Bovens and Hartmann's measure – appears to be inconsistent with their view of coherence as a confidence boosting property.

To see this, consider two independent witnesses with fixed reliability.⁸ Assume furthermore that they give equivalent reports, i.e., that the propositions they report as being true overlap completely. Next, suppose a new witness enters the stage, similar to the first two in that this witness too has a fixed reliability and positively reports on a proposition that is equivalent to the other two. Since the three propositions overlap fully, $a_1 = a_2 = 0$. Also, by theorem 4.1,

$$p^*(R_1 \wedge R_2 \wedge R_3) > p^*(R_1 \wedge R_2) \quad (4.18)$$

⁶One may object here that this example is intuitively not as clear-cut as the other Tweety examples that I have presented. This is no doubt correct. However, for the argument in this section our intuitive judgements do not seem to matter much. Given Bovens and Hartmann's theory of coherence, it simply must be the case that if the posterior probability increases for all values of r , while the prior probability remains the same, then the coherence increases also. But, as we have seen, this is not the case if we measure coherence by $c_r^d(S)$.

Given that $c_r^d(S)$ squares much better with our intuitions in the Tweety example, one may also wonder if $f_r^d(S, S')$ performs better with respect to my other counterexamples against $f_r(S, S')$. This is not the case, however: in all the other counterexamples against Bovens and Hartmann's theory that I have presented in this thesis, $f_r^d(S, S')$ agrees with $f_r(S, S')$ that they constitute indeterminate cases.

⁷Interestingly, the three difference functions discussed so far agree with each other with respect to almost all the examples given in Bovens and Hartmann's (2003a). The only two exceptions are the example they use to criticize Fitelson's measure of coherence and the delta/epsilon version of the corpse in Tokyo example. Both of these will be discussed briefly in section 4.6.

⁸One may assume that the witnesses are actual persons reporting on propositions and that long experience with these people has taught us the true- and false-positive rates of their reports. Alternatively, one may assume the witnesses to be independent measurement instruments of which we likewise know the rates of true- and false-positives.

for all $a_0 \in (0, 1)$: our confidence that the information is true has gone up. This is what we would expect, of course: various independent reports that something is the case will increase our confidence that it is in fact the case. However, in this example all three of the determinants of our confidence in a set remained constant: the reliability was fixed, the prior probability remained equal, and the set was maximally coherent to begin with, so the coherence must have remained maximal also. Therefore, given the fact that according to Bovens and Hartmann there are only three factors that determine our confidence in an information set, we have derived an inconsistency: since all three determinants of our confidence have remained equal, our confidence cannot have increased, while our confidence as measured by $p^*(S)$ did in fact increase.

We have thus arrived at a very general argument against Bovens and Hartmann's theory of coherence, one that is in fact independent of the measure of confirmation with which we measure the probability boost. Therefore, we cannot salvage *Bayesian Coherentism* by replacing the ratio measure in our measure of the confidence boost by another measure of support. Fortunately, I believe that this does not mean that we have derived a general argument against *Bayesian Coherentism*. For there are still a number of options left.

One of these is to allow for a fourth factor that determines our confidence that the information in an information set is true. Evidently, both of my arguments would cease to be compelling if we could find a fourth factor and if it would turn out that this factor does not remain constant if we add a proposition to an information set. For in that case the increase in our confidence that the information is true may well be caused by an increase in the fourth factor and, consequently, the fact that the coherence does not increase need not pose a problem for Bovens and Hartmann's theory of coherence. In their (2003: 21–22), Bovens and Hartmann discuss the possibility of a fourth determinant of our confidence that the information in an information set is true. This question is important for them: clearly, their impossibility result only goes through if we cannot find such a fourth factor. In section 4.2 I agreed with Bovens and Hartmann that it is completely unclear what that factor might be.

However, in Bovens and Hartmann's impossibility result, the number of witnesses is the same for both sets. This clearly is not the case for my argument against their theory. Therefore, could we not simply say that our confidence that the information in an information set is true also depends on the number of witnesses?

I believe we can. Nonetheless, I also believe that this type of solution would present us with a number of difficulties. Firstly, the additional factor is quite *ad hoc* and it seems – at least to me – not as straightforward as the other three factors. Secondly, and perhaps more importantly, a proponent of this solution would also have to provide us with an explanation of why in the original Tweety

example the coherence increases, while in the Tweety example that I presented in the previous section this is no longer the case. Intuitively the cases seem to be very much alike and, consequently, it seems that if the coherence increases in the one case the same should hold for the other one. Here an appeal to specificity along the lines of Bovens and Hartmann (2005) (as discussed in section 3.6) seems to be of little help. For it seems that both sets are equally specific and, therefore, considerations of specificity cannot explain the counterintuitiveness of the results. Thirdly and finally, this approach has the disadvantage that it does not solve any of the counterexamples presented in the previous chapters. That is, according to this theory it will remain the case that in some intuitively very clear cases our theory of coherence as a confidence boosting property will still not judge either of the two sets to be more coherent than the other.

Alternatively, one may attempt to replace the maximality requirement by another condition. In section 4.3, I argued that no measure of coherence can be constructed on the basis of the idea that one should have more confidence in a more coherent set, given equal reliability and equal prior probability. The reason for this is that prior probability and coherence are not independent notions. Of course this need not be an impediment to the construction of a measure of coherence: as intimated above, Bovens and Hartmann simply add a requirement (to wit, the maximality requirement) from which their measure follows quite straightforwardly. Clearly, if we drop this requirement, the argument that Bovens and Hartmann's account is necessarily inconsistent will no longer go through. For in that case the coherence of the set consisting of two equivalent propositions need no longer be maximally coherent, and, therefore, the coherence may increase if we add an additional proposition. Unfortunately, it is unclear what other requirement could take the place of the maximality requirement, so I will not pursue this option any further.

However, I do believe that this section's inconsistency argument does show us that the maximality requirement cannot lead to a satisfactory explication of coherence as a confidence boosting property. Therefore, it needs to be abandoned. However, instead of replacing it by a different condition, I will attempt to slightly alter Bovens and Hartmann's account of what determines our confidence in an information set. It will appear that a relatively minor change will lead to some very different results.

4.6 An Alternative to Bovens and Hartmann's Difference Function

In this section I will outline a different approach to coherence as a confidence boosting property. I will try to stay as close as possible to Bovens and Hartmann's

original account. For example, in my explication of coherence as a confidence boosting property, coherence will still be a property of an information set and it will remain a confidence boosting property in the sense of BC*. What will be different, however, is the explication of the *ceteris paribus* condition in

- BC₁ For all information sets $S, S' \in \mathbf{S}$, if S is no less coherent than S' , then our degree of confidence that the content of S is true is no less than our degree of confidence that the content of S' is true, *ceteris paribus*.

Instead of requiring equal *prior* probability of the conjunction of elements of an information set, I will require equal *unconditional* probabilities of the various elements of an information set. That is, according to me, one should have more confidence in an information set if it is more coherent, given equal reliability and equal *unconditional* probabilities. This will make our confidence in an information set depend on the unconditional probabilities of the various elements of the set, the witness reliability and the coherence, but no longer (directly) on the prior probability of the conjunction of the elements in the set. In that case, a set becomes more coherent if the marginal probabilities and the reliability of the witnesses remain equal while the value of $p^*(R_1 \wedge \dots \wedge R_n)$ increases.

But, rather surprisingly perhaps, it has now become possible to construct a measure on the basis of only this description of coherence as a confidence boosting property. For while it is certainly not the case that prior probability and coherence are independent, it is not evident at all that unconditional probabilities and coherence cannot be modeled as independent. For example, it has become clear in chapter 2 that an explication of coherence as mutual support assumes sets consisting of independent propositions to have neutral coherence, no matter how plausible the propositions are. This at least seems to strongly suggest that the unconditional probabilities do not determine the degree of a set's coherence as mutual support. Naturally, this does not show that the same should apply to an explication of coherence as a confidence boosting property, but it does indicate that it is a quite plausible option.

But if the marginal probabilities of the propositions and the coherence of an information set are indeed independent, then it seems to follow naturally that we should compare the actual probability boost with the probability boost the set would have gotten had all its propositions been independent (but with the same unconditional probabilities and the same witness reliability).

To formalize this idea, again consider the set \mathbf{S} of all information sets. But now define the notion $p_{\text{ind}}^*(R_1 \wedge \dots \wedge R_n)$ as the posterior probability of the information set $S = \{R_1, \dots, R_n\}$, had the information come in as independent propositions with the same unconditional probabilities. Call the set consisting of the same propositions as S , with the same unconditional probabilities, but in which these propositions are independent, set S^i and let $b_{\text{ind}}^i(R_1 \wedge \dots \wedge R_n)$

be the probability boost of the independent set. Next, let a_j^i stand for the sum of the probabilities of the conjunctions consisting of $n - j$ elements of S^i and the negations of all the remaining elements of S^i . From these elements, one can construct a new coherence measure:⁹

$$c_r^i(S) =_{df} \frac{b^i(R_1 \wedge \dots \wedge R_n)}{b_{\text{ind}}^i(R_1 \wedge \dots \wedge R_n)} \quad (4.19)$$

$$= \frac{\frac{p^*(R_1 \wedge \dots \wedge R_n)}{a_0^i}}{\frac{p_{\text{ind}}^*(R_1 \wedge \dots \wedge R_n)}{a_0^i}} \quad (4.20)$$

$$= \frac{p^*(R_1 \wedge \dots \wedge R_n)}{p_{\text{ind}}^*(R_1 \wedge \dots \wedge R_n)} \quad (4.21)$$

$$= \frac{a_0 \sum_{i=0}^n a_i^i (1-r)^i}{a_0^i \sum_{i=0}^n a_i (1-r)^i}, \quad (4.22)$$

and a new difference function:

$$f_r^i(S, S') = c_r^i(S) - c_r^i(S'). \quad (4.23)$$

As in the original case, $S \succsim_i S'$ iff $f_r^i(S, S') \geq 0$ for all values of $r \in (0, 1)$.

Already at first glance, this account has some important advantages over the old one. For one, it now becomes possible to define the properties of coherence and incoherence of a set, which I will call i-(in)coherence:

Definition 4.1 *A set S is i-coherent iff $c_r^i(S) > 0$ for all values of $r \in (0, 1)$; it is i-incoherent iff $c_r^i(S) < 0$ for all values of $r \in (0, 1)$.*

Also, let me stress that this alternative measure leaves intact most of Bovens and Hartmann's model: coherence is still a confidence boosting property of an information set, and the independence of the propositions of the information set is still analyzed in terms of independent witnesses reporting on propositions in an information set.

But more importantly, it can be shown that for all sets S and S' in which the propositions have the same marginal probabilities and for which $p^*(S) > p^*(S')$, it will be the case that S is more i-coherent than S' . To see this, note that if the

⁹One may wonder why I have defined the probability boost as the fraction of the posterior probability and the prior probability the set would have had if all the propositions had been independent. The quick response would be that otherwise it is no longer the case that the resulting difference function is indeterminate in all cases where there is an impossibility result. But I believe it also makes sense to define the probability boost in this fashion, since in this model it is no longer the prior probability that influences our confidence, but the coherence together with the marginal probabilities.

marginal probabilities of the proposition of the two sets are equal it is the case that

$$p_{\text{ind}}^*(S) = p_{\text{ind}}^*(S') \quad (4.24)$$

and thus that

$$f_r^i(S, S') = \frac{\sum_{i=0}^n a_i^i (1-r)^i}{a_0^i} \left(\frac{a_0}{\sum_{i=0}^n a_i (1-r)^i} - \frac{a'_0}{\sum_{i=0}^n a'_i (1-r)^i} \right), \quad (4.25)$$

Evidently, $f_r^i(S, S') = 0$ iff $p^*(S) - p^*(S') = 0$ and $f_r^i(S, S') > 0$ iff $p^*(S) - p^*(S') > 0$, just as we should expect if coherence is to be a confidence boosting property. Personally, I believe that this fact constitutes the main reason why my version of the difference function is a more satisfactory explication of coherence as a confidence boosting property than Bovens and Hartmann's original function. Nonetheless, it also has another important advantage over the original difference function: it appears to square much better with our intuitive notion of coherence.

Firstly, it gives the same answers to most of the examples that Bovens and Hartmann present in their (2003a). To be precise, the new measure disagrees with Bovens and Hartmann's intuitions in only three examples. The first of these is the example that Bovens and Hartmann (2003a: 50–53) use to criticize Fitelson's (2003) measure, which I discussed in section 2.6. I argued that their contention that the coherence of a set should increase if a_0 increases while a_1 is left untouched is based implicitly on intuitions of relative overlap, and thus is not at home in an analysis of coherence as mutual support. Also, I argued in chapter 3 that our intuitions of relative overlap are strongly connected to Fitelson's version of what this chapter has dubbed the maximality requirement. Although there is room for disagreement, I am inclined to argue that by dropping the maximality requirement the justification for Bovens and Hartmann's contention disappears. The same, or so I believe, is the case with the example that Bovens and Hartmann (2003a: 50) use to criticize Shogenji's measure. I will not discuss the example in detail, but let it suffice to say that here, too, an appeal is made to the idea of equivalent propositions being very coherent.

The final example on which the two theories of coherence differ might be considered more important, since it is the only example that Bovens and Hartmann (2003a: 40–41) give of two sets between which no coherence ordering is possible. The example is a corpse in Tokyo example, with a 100-square grid (for a description of the corpse in Tokyo example, see section 3.6). Now consider situations δ and ϵ : in situation δ the two witnesses report squares 41–60 and 51–70. In situation ϵ the witnesses report squares 39–61 and 50–72.

Define ' S^δ ' and ' S^ϵ ' as the information sets containing the witness reports of situations δ and ϵ , respectively. Given the above probabilities, we have $a_0^\delta = .10$, $a_1^\delta = .20$, $a_0^\epsilon = .12$, and $a_1^\epsilon = .22$. Since $f(S^\delta, S^\epsilon)$ is positive for values of r between 0 and, approximately, .7 and negative for values of r between, approximately, .7

and 1, we have neither $S^\delta \succcurlyeq S^\epsilon$ nor $S^\epsilon \succcurlyeq S^\delta$. However, as figure 4.4 shows, $f_r^i(S^\delta, S^\epsilon)$ is larger than 0 for all values of r , and therefore, $S^\delta \succcurlyeq_i S^\epsilon$.




Figure 4.4: Graph of $f_r^i(S^\delta, S^\epsilon)$

According to the difference function there is no fact of the matter as to which set is more coherent, and Bovens and Hartmann believe that this constitutes a point in favor of their theory. In such cases, they think, a theory of coherence should withhold judgement. But it seems telling that almost all the other recently proposed probabilistic measures of coherence side with $f_r^i(S^\delta, S^\epsilon)$. Fitelson's, Shogenji's and all of the mutual support measures proposed in chapter 2 all point to S^δ as being more coherent than S^ϵ .¹⁰ Only according to Olsson's (2002) measure and the overlap measure of coherence \mathcal{R} proposed in chapter 3 is the opposite the case.¹¹

So there is at least some reason to argue that S^δ is more coherent than S^ϵ . Moreover, this conclusion is not incompatible with Bovens and Hartmann's claim that in this case it is intuitively not clear which set is more coherent. We may intuitively be unable to see the difference in coherence for the simple reason that it is too small for us to recognize.

As can be seen from the values of the respective measures, according to all of the measures mentioned above the differences in coherence between the two sets are indeed minute, which would explain our being unable to identify the most coherent set. Of course, Bovens and Hartmann's framework does not supply a mechanism with which to decide *how much* more coherent a set is than another, which makes it impossible to distinguish between intuitively very clear cases and intuitively less clear cases. This is just a natural result of their approach.

¹⁰Shogenji's (1999) measure yields $S^\delta = 2.5$ and $S^\epsilon \approx 2.27$; Fitelson's (2003) measure yields $S^\delta = .6$ and $S^\epsilon \approx .57$, the three measures discussed in chapter 2 yield $C_d(S^\delta) = .3$, $C_d(S^\epsilon) \approx .29$, $C_r(S^\delta) = 2.5$, $C_r(S^\epsilon) \approx 2.27$, $C_l(S^\delta) = 4.0$, $C_l(S^\epsilon) \approx 3.6$.

¹¹Both measures yield $S^\delta \approx .333$ and $S^\epsilon \approx .353$.

However, it is quite easy to add such an element to $f_r^i(S, S')$. For example, one could take the maximum value of $f_r^i(S, S')$ or the area beneath its graph to be indicative of the degree to which S is more coherent than S' (provided $f_r^i(S, S') \geq 0$ for all r). In the case of S^δ and S^ϵ , both proposals would lead to the conclusion that the difference in coherence between S^δ and S^ϵ is much smaller than in all of the other examples discussed in this thesis.

Therefore, I am inclined to conclude that my version of the difference function gives the correct judgements in all the intuitively clear examples presented by Bovens and Hartmann. But be that as it may, it certainly behaves much more in line with our coherence intuitions with respect to the counterexamples that this thesis has presented against their theory of coherence. Indeed, it gives the correct answer to all of the counterexamples against the original difference function, with only one exception. In this example (which I presented in subsection 2.5.4) we consider a set $S = \{C, E, T\}$ with

- C: This chair is brown;
- E: Electrons are negatively charged;
- T: Today is Thursday.

and compare it with another set $S' = \{B, O, M\}$ with

- B: This bird is black;
- O: This bird is a crow;
- M: This bird has a lifelong mate.

Now suppose the probabilities for the two sets are given by diagram 4.5 (this is the counterexample against Bovens and Hartmann's difference function given the *ceteris paribus* clause that the prior probabilities remain equal). Presumably, set S' is intuitively much more coherent than set S . However, according to my version of the difference function, it is indeterminate whether S' is more i -coherent than set S . Of course, this example does not constitute a point in favor of Bovens and Hartmann's theory, since their difference function likewise does not classify one of the sets as more coherent. Nevertheless, one may feel that we do have a counterexample against my own difference function. But do we?

I believe that whether or not the above example constitutes a counterexample against a theory of coherence depends on the notion of coherence we attempt to explicate. If we are after a notion of coherence as mutual support, relative overlap, or a combination of the two, then the example constitutes a true counterexample. For it is both the case that the mutual support is higher in the second set than in the first and that the relative overlap is larger in the second set than in the first. However, if we are after an explication of the concept of coherence as a confidence boosting property, I am not so sure what the intuitively correct judgement should be. For example, one plausible explanation why S' should be more coherent than S is that the first set is independent, while in the other set all propositions support




Figure 4.5: New diagrams of the probability distributions corresponding to sets *S* (left) and *S'* (right)

each other. But we need to remember that we are trying to explicate the notion of coherence as a confidence boosting property and it need not be the case that the posterior probability of a positively dependent set is higher for all values of r than that of an independent set. Indeed, precisely the two sets above constitute a counterexample against that claim. For consider figure 4.6, which gives the difference in posterior probability. Clearly, this graph crosses the r -axis.




Figure 4.6: Graph of $p^*(S') - p^*(S)$

Thus, it appears that given quite reliable witnesses an any-any coherent set (as defined by definition 2.4) can have a lower posterior probability than an independent set and, furthermore, a set with a higher relative overlap can likewise be less probable than a set with a lower relative overlap. This remark is not meant

to imply that the example *should* constitute an indeterminate case (for the marginal probabilities are not equal in both sets), but it does at least suggest that our intuitions with respect to this example may well be less convincing from the standpoint of the explication of coherence as a confidence boosting property.

As a last remark, let me note that although this measure does not satisfy the maximality requirement, it does satisfy a less stringent condition, to wit, that sets consisting of equivalent propositions are maximally coherent given constant unconditional probabilities and equal reliability of the witnesses (for a proof, see Appendix 4 C):

Theorem 4.3 *Any equivalent set $S = \{R_1, \dots, R_n\}$ with unconditional probabilities $p(R_i) = \alpha$ is more i-coherent than any nonequivalent set S' with the same number of propositions and the same unconditional probabilities.*

Thus, given two sets with the same unconditional probabilities for all the propositions, if only one of them is equivalent it will be more i-coherent than the other. Or, alternatively, given fixed unconditional probabilities, the most i-coherent way information can accumulate is if the propositions are all equivalent.

4.7 Conclusion

In this chapter I have explicated the notion of coherence as a confidence boosting property. From Bovens and Hartmann's impossibility result it follows that coherence can only impose a quasi-ordering on the set of information sets. But whereas Bovens and Hartmann's difference function does indeed impose such a quasi-ordering, it appeared that it also refrains from imposing an ordering in cases where such an ordering is required.

I have argued that there are only two viable solutions to this problem, both of which make some alterations in Bovens and Hartmann's model for the factors that determine our confidence in an information set given positive reports by all of the witnesses. The first solution is simply to add an additional factor, namely, the number of witnesses. Although this solves the counterarguments presented in this chapter, I find this solution less satisfactory than the second one, in which our confidence in an information set is still determined jointly by the probability of the results, the reliability of the witnesses and the coherence of the results, but in which we explicate the probability of the results in terms of the marginal probabilities of the respective propositions. The latter solution has the advantage of satisfactorily dealing with both the problems discussed in this chapter and those presented in the other chapters.

Let me emphasize that I have discussed only one alternative to Bovens and Hartmann's original difference function. Most notably, I have followed Bovens

and Hartmann in defining the probability boost via the ratio measure. But I could just as well have defined the probability boost by means of the difference measure or the likelihood measure (or any other measure of confirmation, for that matter). Alternatively, I could also have defined the coherence of a set as the difference between the posterior probability and the posterior probability the set would have had if the propositions had been independent. This would result in the following coherence function:

$$c'_r =_{df} p^*(S) - p^*_{\text{ind}}(S) \quad (4.26)$$

None of these alternatives are ordinally equivalent. For example, replacing the ratio-measure by the difference measure in the definition of $c^i_r(S)$ will lead to the delta/epsilon version of the Tokyo example constituting an indeterminate case. However, the sets presented in Bovens and Hartmann's impossibility result (see section 4.2) would cease to constitute an indeterminate case. Alternatively, using expression (4.26) for the definition of the coherence function seems to lead much more often to an indeterminate judgement (but not in any of the counterexamples that this thesis has brought forward against Bovens and Hartmann's difference function).

So far I have not succeeded in finding a definite 'test-case' against any of these alternative formulations and for the sake of clarity I have opted in favor of only discussing the alternative that comes closest to Bovens and Hartmann's original formulation. Nonetheless, if somebody were to produce a decisive counterexample against the measure of i-coherence proposed in this chapter, it would be interesting to find out whether all of the other options lead to the same results.

Appendix 4 A: Proofs of Theorems 4.1 and 4.2

Theorem 4.1 Given two sets $S = \{R_1, \dots, R_{n-1}\}$ and $S' = \{R_1, \dots, R_n\}$, and with $0 < p(R_n) < 1$ and $p(R_1 \wedge \dots \wedge R_{n-1}) = p(R_1 \wedge \dots \wedge R_n)$, it will be the case that $p^*(R_1 \wedge \dots \wedge R_n) > p^*(R_1 \wedge \dots \wedge R_{n-1})$ for all $r \in (0, 1)$.

Proof: Consider two sets $S = \{R_1, \dots, R_{n-1}\}$ and $S' = \{R_1, \dots, R_n\}$. For convenience, substitute b_i 's for the a_i 's of S' . Assume that $p(R_1 \wedge \dots \wedge R_{n-1}) = p(R_1 \wedge \dots \wedge R_n)$. Then, $a_0 = b_0$. Next consider the difference in posterior probabilities of the two sets:

$$\begin{aligned} & p^*(R_1 \wedge \dots \wedge R_n) - p^*(R_1 \wedge \dots \wedge R_{n-1}) \\ &= \frac{b_0}{\sum_{i=0}^n b_i (1-r)^i} - \frac{a_0}{\sum_{i=0}^{n-1} a_i (1-r)^i} \\ &= a_0 \left(\frac{1}{\sum_{i=0}^n b_i (1-r)^i} - \frac{1}{\sum_{i=0}^{n-1} a_i (1-r)^i} \right), \end{aligned} \quad (4.27)$$

where I have made use of the fact that $a_0 = b_0$. In order to prove the theorem, I will have to show that if $0 < p(R_n) < 1$, equation (4.27) is positive for all values of r . For this to hold, it must be the case that

$$\sum_{i=0}^{n-1} a_i (1-r)^i - \sum_{i=0}^n b_i (1-r)^i \quad (4.28)$$

is positive for all values of r . In order to compare the two sums above, it is necessary to derive a relation between the a_i 's and the b_i 's. Let us start with dividing the b_i 's into two parts: $b_i^{R_n}$ is the part of b_i where R_n is true and $b_i^{\neg R_n}$ is the part of b_i where R_n is false. In that case, $b_i^{R_n} + b_i^{\neg R_n} = b_i$ and $b_0^{\neg R_n} = b_n^{R_n} = 0$. Also:

$$\sum_{i=0}^n b_i^{\neg R_n} = \sum_{i=1}^n b_i^{\neg R_n} = 1 - p(R_n). \quad (4.29)$$

Furthermore, b_0 can be expressed in terms of a_0 and $b_1^{\neg R_n}$ by using the law of total probability:

$$\begin{aligned} b_0 &= p(R_1 \wedge \dots \wedge R_{n-1} \wedge R_n) \\ &= p(R_1 \wedge \dots \wedge R_{n-1}) - p(R_1 \wedge \dots \wedge R_{n-1} \wedge \neg R_n) \\ &= a_0 - b_1^{\neg R_n}. \end{aligned} \quad (4.30)$$

The same type of relation can be derived for all b_i . By definition, $b_i^{R_n}$ is the sum of all possible combinations in which R_n is true and precisely i propositions are

false:

$$b_i^{R_n} = \sum_{\substack{i\text{-many } \neg R_k's, \\ (n-1-i)\text{-many } R_k's}} p(R_1 \wedge \cdots \wedge R_{n-1} \wedge R_n),$$

with $k \leq n-1$. Again by the law of total probability,

$$\begin{aligned} b_i^{R_n} &= \sum_{\substack{i\text{-many } \neg R_k's, \\ (n-1-i)\text{-many } R_k's}} (p(R_1 \wedge \cdots \wedge R_{n-1}) - p(R_1 \wedge \cdots \wedge R_{n-1} \wedge \neg R_n)) \\ &= \sum_{\substack{i\text{-many } \neg R_k's, \\ (n-1-i)\text{-many } R_k's}} p(R_1 \wedge \cdots \wedge R_{n-1}) - \sum_{\substack{i\text{-many } \neg R_k's, \\ (n-1-i)\text{-many } R_k's}} p(R_1 \wedge \cdots \wedge R_{n-1} \wedge \neg R_n). \end{aligned}$$

The first term on the right-hand side of the last equation equals a_i : it is the sum of all the joint probabilities of all combinations of set S where i propositions are false and all others are true. The second term equals $b_{i+1}^{\neg R_n}$: it is the sum of the joint probabilities of all combinations of set S' in which R_n and precisely i other propositions are false. Therefore:

$$b_i^{R_n} = a_i - b_{i+1}^{\neg R_n}$$

and

$$b_i = b_i^{R_n} + b_i^{\neg R_n} = a_i - b_{i+1}^{\neg R_n} + b_i^{\neg R_n}, \quad (4.31)$$

where $a_n = b_{n+1} = 0$. Substituting expression (4.31) in equation (4.28) yields:

$$\begin{aligned} &\sum_{i=0}^{n-1} a_i (1-r)^i - \sum_{i=0}^n b_i (1-r)^i \\ &= \sum_{i=0}^{n-1} a_i (1-r)^i - \sum_{i=0}^n (a_i - b_{i+1}^{\neg R_n} + b_i^{\neg R_n}) (1-r)^i \\ &= \sum_{i=0}^{n-1} a_i (1-r)^i - \sum_{i=0}^n a_i (1-r)^i + \sum_{i=0}^n (b_{i+1}^{\neg R_n} - b_i^{\neg R_n}) (1-r)^i \\ &= \sum_{i=0}^{n-1} b_{i+1}^{\neg R_n} (1-r)^i - \sum_{i=1}^n b_i^{\neg R_n} (1-r)^i, \end{aligned} \quad (4.32)$$

where I have made use of the fact that $a_n = b_{n+1} = b_0^{\neg R_n} = 0$. Equation (4.32) can be further simplified:

$$\sum_{i=0}^{n-1} b_{i+1}^{\neg R_n} (1-r)^i - \sum_{i=1}^n b_i^{\neg R_n} (1-r)^i$$

$$\begin{aligned}
&= \sum_{i=1}^n b_i^{-R_n} (1-r)^{i-1} - \sum_{i=1}^n b_i^{-R_n} (1-r)^i \\
&= \sum_{i=1}^n \left(b_i^{-R_n} (1-r)^{i-1} - b_i^{-R_n} (1-r)^i \right) \\
&= \sum_{i=1}^n b_i^{-R_n} (1-r)^{i-1} (1 - (1-r)) \\
&= \sum_{i=1}^n b_i^{-R_n} r (1-r)^{i-1}. \tag{4.33}
\end{aligned}$$

Suppose $p(R_n) < 1$. Then equation (4.33) is positive for all values of r and, therefore, $p^*(R_1 \wedge \cdots \wedge R_n) > p^*(R_1 \wedge \cdots \wedge R_{n-1})$ for all $r \in (0, 1)$. ■

Theorem 4.2 *Adding a proposition R_{m+1} to a nonequivalent set $S = \{R_1, \dots, R_m\}$ will increase S 's coherence iff R_{m+1} is equivalent with $\bigwedge S$, i.e., iff $p(R_1 \wedge \cdots \wedge R_m \mid R_{m+1}) = p(R_{m+1} \mid R_1 \wedge \cdots \wedge R_m) = 1$.*

Proof: I will start by showing that adding a proposition R_{m+1} to any set $S = \{R_1, \dots, R_m\}$ will not increase S 's coherence if R_{m+1} is not equivalent with S . To this end, consider two sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R_1, \dots, R_m, R_n\}$, where $n = m + 1$. Again, substitute b_i 's for the a_i 's of S' .

Either $b_0 = 0$ or $b_0 > 0$. In the first case, $f(S', S)$ is not defined, and, therefore, S' cannot be more coherent than S . Thus, b_0 must be larger than 0 (and, consequently, $a_0 > 0$ also).

Note that it would suffice for S' to be no more coherent than S if the derivative of $f_r(S', S)$ is negative in $r = 0$. Using equation (3.16), the derivative of $f(S', S)$ for sets $S = \{R_1, \dots, R_m\}$ and $S' = \{R_1, \dots, R_m, R_n\}$ equals

$$\begin{aligned}
\left(\frac{df(S', S)}{dr} \right)_{r=0} &= \sum_{i=1}^{n-1} (i-n)b_i - \sum_{i=1}^{m-1} (i-m)a_i \\
&= \sum_{i=1}^{m-1} (m-i)a_i - \sum_{i=1}^{n-1} (n-i)b_i. \tag{4.34}
\end{aligned}$$

Substituting expression (4.31) in expression (4.34) and using $n = m + 1$ gives:

$$\begin{aligned}
&\sum_{i=1}^{m-1} (m-i)a_i - \sum_{i=1}^{n-1} (n-i)b_i \\
&= \sum_{i=1}^{m-1} (m-i)a_i - \sum_{i=1}^{n-1} (n-i) \left(a_i + b_i^{-R_n} - b_{i+1}^{-R_n} \right)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^{m-1} (m-i)a_i - \sum_{i=1}^m (m+1-i)a_i - \sum_{i=1}^m (n-i)b_i^{-R_n} + \sum_{i=1}^m (n-i)b_{i+1}^{-R_n} \\
&= \sum_{i=1}^{m-1} ((m-i)a_i - (m+1-i)a_i) - a_m - \sum_{i=1}^m (n-i)b_i^{-R_n} + \sum_{i=2}^{m+1} (n-i+1)b_i^{-R_n} \\
&= \sum_{i=1}^{m-1} (-a_i) - a_m - \sum_{i=1}^m (n-i)b_i^{-R_n} + \sum_{i=2}^{m+1} (n-i+1)b_i^{-R_n} \\
&= \sum_{i=1}^m (-a_i) - (n-1)b_1^{-R_n} - \sum_{i=2}^m (n-i)b_i^{-R_n} + \sum_{i=2}^m (n-i+1)b_i^{-R_n} + b_n^{-R_n} \\
&= -\sum_{i=1}^m a_i - mb_1^{-R_n} + \sum_{i=2}^m b_i^{-R_n} + b_n^{-R_n} \\
&= -\sum_{i=1}^m a_i + \sum_{i=1}^n b_i^{-R_n} - b_1^{-R_n} - mb_1^{-R_n}.
\end{aligned}$$

Using equations (4.29) and (4.30) and substituting $\sum_{i=1}^m a_i = 1 - a_0$, it is possible to derive a simple expression for the derivative of $f(S', S)$ for $r = 0$:

$$\begin{aligned}
\left(\frac{df(S', S)}{dr} \right)_{r=0} &= -1 + b_0 + b_1^{-R_n} + 1 - p(R_n) - b_1^{-R_n} - mb_1^{-R_n} \\
&= b_0 - p(R_n) - mb_1^{-R_n} \\
&= p(R_1 \wedge \dots \wedge R_m \wedge R_n) - p(R_n) - \\
&\quad mp(R_1 \wedge \dots \wedge R_m \wedge \neg R_n) \\
&= p(R_n) (p(R_1 \wedge \dots \wedge R_m | R_n) - 1) - \\
&\quad mp(\neg R_n) p(R_1 \wedge \dots \wedge R_m | \neg R_n).
\end{aligned}$$

From this it follows directly that the derivative of $f(S', S)$ for $r = 0$ is negative if $p(R_1 \wedge \dots \wedge R_m | R_n) < 1$, or if $p(R_1 \wedge \dots \wedge R_m | \neg R_n) > 0$. Therefore, adding a proposition R_{m+1} to a set of propositions $S = \{R_1, \dots, R_m\}$ will not increase its coherence if R_{m+1} is not equivalent with S .

If R_{m+1} is equivalent with S , coherence will increase if set S is nonequivalent. For consider the coherence function for set $S' = \{R_1, \dots, R_m, R_n\}$:

$$c_r(S') = \frac{b_0 + (1 - b_0)(1 - r)^n}{\sum_{i=0}^n b_i (1 - r)^i}.$$

If R_n is equivalent with $R_1 \wedge \dots \wedge R_m$, then $b_0 = a_0$, $b_1 = 0$ and $b_{i+1} = a_i$ for $i \geq 1$. Therefore:

$$c_r(S') = \frac{a_0 + (1 - a_0)(1 - r)^n}{a_0 + \sum_{i=2}^n a_{i-1}(1 - r)^i}$$

$$\begin{aligned}
&= \frac{a_0 + (1 - a_0)(1 - r)^n}{a_0 + \sum_{i=1}^m a_i(1 - r)^{i+1}} \\
&= \frac{\frac{a_0}{(1-r)} + (1 - a_0)(1 - r)^m}{\frac{a_0}{(1-r)} + \sum_{i=1}^m a_i(1 - r)^i}.
\end{aligned}$$

It is fairly easy to show that for all $\alpha, \beta, \gamma, \delta \in (0, 1)$, if $\gamma < \delta$,

$$\frac{\frac{\alpha}{\beta} + \gamma}{\frac{\alpha}{\beta} + \delta} > \frac{\alpha + \gamma}{\alpha + \delta}.$$

If S is nonequivalent, then

$$\sum_{i=1}^m a_i(1 - r)^m < \sum_{i=1}^m a_i(1 - r)^i,$$

and thus $(1 - a_0)(1 - r)^m < \sum_{i=1}^m a_i(1 - r)^i$. Therefore, for all $0 < r < 1$ and all $0 < a_0 < 1$, we have

$$\begin{aligned}
c_r(S') &> \frac{a_0 + (1 - a_0)(1 - r)^m}{a_0 + \sum_{i=1}^m a_i(1 - r)^i} \\
&= c_r(S).
\end{aligned}$$

Therefore, $f_r(S', S) =_{df} c_r(S') - c_r(S) > 0$ for all values of r and thus S' is more coherent than S . ■

Appendix 4 B: Derivation of the Measure $c_r^l(S)$

Consider a set $S = \{R_1, \dots, R_n\}$. Measuring the probability boost by means of the likelihood measure gives the following expression:

$$b^l(R_1 \wedge \dots \wedge R_n) = \frac{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n)}{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | \neg(R_1 \wedge \dots \wedge R_n))}.$$

First use Bayes's theorem to obtain:

$$\begin{aligned} & \frac{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n)}{p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) \frac{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}{p(\neg(R_1 \wedge \dots \wedge R_n))}} \\ = & \frac{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n) p(\neg(R_1 \wedge \dots \wedge R_n))}{p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} \\ = & \frac{p(\text{REPR}_1 | R_1) \dots p(\text{REPR}_n | R_n) p(\neg(R_1 \wedge \dots \wedge R_n))}{p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} \\ = & \frac{p^n(1 - a_0)}{p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}, \end{aligned}$$

where p is the true-positive rate of the witnesses ($p =_{df} p(\text{REPR}_j | R_j)$, see also section 4.2) and where I have made use of the fact that $p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n) = p(\text{REPR}_1 | R_1) \dots p(\text{REPR}_n | R_n)$, which is proven by Bovens and Hartmann (2003a: 133). Also by Bayes's theorem, we have:

$$\begin{aligned} & \frac{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n | R_1 \wedge \dots \wedge R_n) p(R_1 \wedge \dots \wedge R_n)} \\ = & \frac{p^n a_0}{p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}. \end{aligned}$$

and thus,

$$p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) = \frac{p^n a_0}{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}.$$

Filling this in in the expression above and using the negation rule gives:

$$\begin{aligned} & \frac{p^n(1 - a_0)}{p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) p(\text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} \\ = & \frac{(1 - a_0) p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}{a_0 p(\neg(R_1 \wedge \dots \wedge R_n) | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} \\ = & \frac{(1 - a_0) p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)}{a_0(1 - p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n))} \end{aligned}$$

$$= \frac{(1 - a_0)}{a_0 \left(\frac{1}{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} - 1 \right)}.$$

Substituting $p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n) = \frac{(1-a_0)}{\sum_{i=1}^n a_i(1-r)^i}$ gives

$$\begin{aligned} & \frac{(1 - a_0)}{a_0 \left(\frac{1}{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \dots \wedge \text{REPR}_n)} - 1 \right)} \\ &= \frac{(1 - a_0)}{a_0 \left(\frac{\sum_{i=1}^n a_i(1-r)^i}{a_0} - 1 \right)} \\ &= \frac{(1 - a_0)}{\sum_{i=0}^n a_i(1-r)^i - a_0} \\ &= \frac{(1 - a_0)}{\sum_{i=1}^n a_i(1-r)^i}. \end{aligned}$$

In case S is an equivalent set, this reduces to:

$$b_{\max}^l(R_1 \wedge \dots \wedge R_n) = \frac{(1 - a_0)}{(1 - a_0)(1 - r)^n} = \frac{1}{(1 - r)^n}.$$

Combining the two gives:

$$\begin{aligned} c_r^l(S) &=_{df} \frac{b^l(R_1 \wedge \dots \wedge R_n)}{b_{\max}^l(R_1 \wedge \dots \wedge R_n)} \\ &= \frac{\frac{(1-a_0)}{\sum_{i=1}^n a_i(1-r)^i}}{\frac{1}{(1-r)^n}} \\ &= \frac{(1 - a_0)(1 - r)^n}{\sum_{i=1}^n a_i(1 - r)^i}. \end{aligned}$$

■

Appendix 4 C: Proof of Theorem 4.3

Theorem 4.3 *Any equivalent set $S = \{R_1, \dots, R_n\}$ with unconditional probabilities $p(R_i) = \alpha$ is more i-coherent than any nonequivalent set S' with the same number of propositions and the same unconditional probabilities.*

Proof: Consider a set $S = \{R_1, \dots, R_n\}$ consisting only of equivalent propositions with probability $p(R_i) = \alpha$. Next consider a second set $S' = \{R'_1, \dots, R'_n\}$ with the same number of propositions, each of which has unconditional probability $p(R'_i) = \alpha$, but in which at least one proposition or subset of propositions is not equivalent to all of the others.

First, rewrite the expression for $p^*(S)$:

$$\begin{aligned} p^*(S) &= \frac{a_0}{\sum_{i=1}^n a_i (1-r)^i} \\ &= \frac{a_0}{a_0 + \sum_{i=1}^{n-1} a_i (1-r)^i + a_n (1-r)^n} \\ &= \frac{1}{1 + \sum_{i=1}^{n-1} \frac{a_i}{a_0} (1-r)^i + \frac{a_n}{a_0} (1-r)^n}. \end{aligned}$$

Evidently, if set S is equivalent $\sum_{i=1}^{n-1} a_i (1-r)^i = 0$, and if set S' is nonequivalent $\sum_{i=1}^{n-1} a'_i (1-r)^i > 0$. Moreover, if the propositions in both sets have the same marginal probabilities, then $a'_0 < a_0$ and $a'_n > a_n$ and therefore, $\frac{a'_n}{a'_0} > \frac{a_n}{a_0}$.

Next, note that since all propositions in both sets have the same unconditional probabilities, the posterior probabilities of the independent sets are equal: $p_{\text{ind}}^*(S) = p_{\text{ind}}^*(S')$. Therefore $f_r^i(S, S') \geq 0$ for all r iff $p^*(S) - p^*(S') > 0$, or, alternatively, if $\frac{1}{p^*(S)} < \frac{1}{p^*(S')}$. As the following shows, this is indeed the case:

$$\begin{aligned} &\frac{1}{p^*(S)} - \frac{1}{p^*(S')} \\ &= \frac{\sum_{i=0}^n a_i (1-r)^i}{a_0} - \frac{\sum_{i=0}^n a'_i (1-r)^i}{a'_0} \\ &= 1 + \frac{a_n}{a_0} (1-r)^n - \left(1 + \sum_{i=1}^{n-1} \frac{a'_i}{a'_0} (1-r)^i + \frac{a'_n}{a'_0} (1-r)^n \right) \\ &= \frac{a_n}{a_0} (1-r)^n - \frac{a'_n}{a'_0} (1-r)^n - \sum_{i=1}^{n-1} \frac{a'_i}{a'_0} (1-r)^i \\ &= \left(\frac{a_n}{a_0} - \frac{a'_n}{a'_0} \right) (1-r)^n - \sum_{i=1}^{n-1} \frac{a'_i}{a'_0} (1-r)^i \\ &< 0. \end{aligned}$$

Since the above result holds for all $r \in (0, 1)$, it follows that S is more i-coherent than S' . ■

Chapter 5

A Probabilistic Measure of Bootstrap Confirmation

5.1 Introduction

It is widely acknowledged that empirical testing crucially involves the use of auxiliary theories and is thus in an uncontroversial sense always relative to some theory or theories. Nonetheless, in his *Theory and Evidence* Glymour has presented an important new confirmation theory according to which empirical theories can be confirmed in an absolute sense nonetheless. Glymour's theory of bootstrap confirmation – as he called it – is a purely qualitative confirmation theory; it allows us to say that the evidence confirms a given theory, but not that it confirms the theory to a certain degree. In the present chapter I will try to take some first steps toward extending Glymour's theory into a quantitative account.

This chapter is organized as follows. In section 5.2, Glymour's theory of bootstrap confirmation is briefly discussed and a qualitative but probabilistic definition of bootstrap confirmation is given. Next, section 5.3 will propose a number of desiderata for a quantitative measure of bootstrap confirmation and present a family of measures all of which satisfy the desiderata. In section 5.4, I will compare the different measures. Although it is clear that they behave strikingly different in a number of cases, I will not make an argument in favor of any of the measures for the simple reason that it seems intuitively unclear what would be the most satisfactory behavior for a measure of bootstrap confirmation in these cases. Next, section 5.5 will discuss a few objections that can be raised against the measures of bootstrap confirmation that the previous sections have proposed.

As first sight it may seem that this chapter is of a very different nature than

the rest of this thesis. But although it is true that this chapter proposes a measure of confirmation instead of coherence, it will become clear that this measure bears some striking resemblances to the measures of coherence proposed above, especially to the class of measures of coherence as mutual support (see chapter 2). Indeed, as I will show in section 5.6, the measure of bootstrap confirmation proposed in this chapter seems to be quite successful at explicating a notion of coherence that so far has not yet been given a formal explication, to wit, that of the coherence between a theory and the evidence. Although, as I will argue there, some alterations in the definition of the measure may seem appropriate if it is to be used for that purpose, it seems that the general framework discussed in this chapter goes at least some way toward explicating this notion of coherence.

5.2 Glymour's Qualitative Theory of Bootstrap Confirmation

A confirmation theory is, roughly put, a theory that purports to specify, for any given evidence statement and any given hypothesis, whether or not the former *supports* the latter, or – in different terms that for present purposes can all be taken as equivalents – whether coming to know the evidence statement should *increase our confidence* in the hypothesis, whether the evidence *adds to the justificational status of the hypothesis*, whether it *gives reason to believe* the hypothesis. A confirmation theory may or may not also specify *to what extent* the evidence supports, or should affect our confidence in, a given hypothesis, *how much* it adds to the justificational status of a hypothesis, *how much* reason it gives to believe the hypothesis. If it does specify the extent of support, it is called a *quantitative* confirmation theory, if it does not, it is called a *qualitative* confirmation theory.

All these formulations suggest that confirmation is a two-place relation, viz., a relation between a body of evidence and a hypothesis. And indeed this is what philosophers for a long time believed. However, as Duhem (1906/1954) was the first to argue, and as Quine (1953) famously repeated in his assault on the logical empiricists' reductionist semantics, confirmation is a three- rather than a two-place relation: evidence generally accrues to a hypothesis only relative to one or more auxiliary hypotheses. It is no exaggeration to say that today this is something of a commonplace among analytic philosophers.¹

¹Subjective Bayesians may want to deny this. On their account, scientists are free – within the bounds of probability theory – in the probabilities they assign, and thus may also assign a probability to a hypothesis conditional upon the evidence alone (not conjoined to any auxiliaries, that is) that is greater than the unconditional probability assigned to the hypothesis (in which case the evidence confirms the hypothesis). However, the quantitative theory of bootstrap confirmation

Many have taken Duhem's thesis to imply that all (dis)confirmation must be 'relative' (dis)confirmation and it seems only one step further to argue for a radical holism according to which empirical claims face experience collectively, not individually. As is generally known, Quine has taken this idea to the extreme and argued that 'any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system' (Quine 1953: 43). Therefore, even the laws of logic and mathematics may be discarded, which, however, would be very inconvenient and therefore not be pragmatically justified (see, for instance, Ladyman 2002: 172).

Glymour is very unhappy with such radically holistic conclusions. For him, the practice of the natural sciences does not warrant any extremely holistic conclusions at all:

Taken literally, a radical holism makes it impossible to understand either the most elaborate or the most trivial of scientific arguments. What Kepler, Galileo, Newton, Dalton, Maxwell, Perrin, Freud, Einstein, Eddington, whoever, were doing in much of their scientific writing – wherever they purported to relate particular pieces of evidence to particular pieces of theory – becomes an utter mystery. Whatever merit there may be in holism suitably qualified, to embrace a radical holism is only to confess that one does not know at all how science succeeds in doing what it plainly does. (Glymour 1980a: 45)

In his eyes, scientists very often relate a case of disconfirmation to a specific part of a theory and are often remarkably capable of altering precisely that part of the theory that seems to cause the trouble. Even when scientists acknowledge the fundamental role auxiliaries play in a certain test, they sometimes still draw the – according to Glymour – valid conclusion that a hypothesis is confirmed or disconfirmed by that test.

In order to justify these ideas, Glymour tries to incorporate the claims by Duhem and Quine that confirmation is relative to auxiliaries while evading the radically holistic conclusions that many have believed follow from accepting these claims. More concretely, he argues that the indispensability of auxiliaries in the testing of *single* hypotheses is no impediment to absolute confirmation of *complexes* of such hypotheses or, as one may call them, theories. To make this general idea precise, Glymour presents a confirmation theory according to which the piecemeal confirmation of the individual hypotheses comprised by a given theory relative to other hypotheses comprised by the same theory *may* add up to an unrelativized confirmation of that theory as a whole. Whether it does depends

to be developed below is neutral on the indispensability of auxiliaries in the sense that the measure of bootstrap support to be proposed also takes into account any support the evidence might give a hypothesis in isolation (in addition to the support the evidence may give the hypothesis relative to various auxiliaries).

on whether the separate tests of the various hypotheses *could* have turned out negative for these hypotheses. That is, only in cases where the auxiliaries do not shield the hypothesis from disconfirmation can true confirmation be achieved.

Since its publication, it has become abundantly clear that Glymour's theory of bootstrap confirmation as it is presented in Glymour (1980a and 1983) and in Earman and Glymour (1988) faces some important problems. Edidin (1983 and 1988), van Fraassen (1983a), Horwich (1980) and Christensen (1983 and 1990) have all shown that the approach suffers from several serious defects. And although Glymour has successfully dealt with some of these challenges, many others remain.² Owing to these problems many have concluded that Glymour's theory is untenable and it seems no exaggeration to say that the theory has practically gone into oblivion (cf. Douven 2005b).

Interestingly, almost all the criticisms of bootstrap confirmation have been directed against Glymour's use of the Hempelian positive instance account of confirmation. Since Glymour has emphasized in his (1980a: 127) and again in his (1980b) that his account of bootstrap confirmation is independent of the specific theory of confirmation that one favors, it is surprising that virtually everybody working on the subject has taken the criticisms of Glymour's Hempelian version of bootstrap confirmation to be indicative of a failure of bootstrap confirmation *per se*. One might have expected – especially due to the potential of bootstrap confirmation to overcome radical holism – the failures of Hempelian bootstrap confirmation to have encouraged others to pursue different accounts of bootstrap confirmation. However, it seems that no such attempts have been made.

One of the reasons for this is undoubtedly that Glymour's (1980a: 130–131) definition of bootstrap confirmation does not distinguish between his bootstrap intuitions and the account of confirmation used to cash out these intuitions. To see whether Glymour's account can be salvaged after all, we therefore require a definition in which these two elements of Glymour's account are separated more clearly than in Glymour's original definition. As a definition of bootstrap confirmation that is independent of the account of confirmation that one favors, I propose the following:³

²Noteworthy in this respect are Christensen's examples of counterintuitive cases of confirmation condoned by Glymour's theory of bootstrap confirmation. One of Glymour's main reasons for rejecting both hypothetico-deductivism (H-D) and Hempel's (1965) account of instance confirmation is that they allow counterintuitive cases of confirmation like the proposition that the moon is made of green cheese confirming the proposition that salt dissolves in water (H-D) and a white shoe confirming the proposition that all ravens are black (Hempel's account). But as Christensen (1983 and 1990) shows, there exist similar examples with respect to Glymour's theory of bootstrap confirmation. For example, according to Glymour's theory a black raven with wings bootstrap confirms the proposition that only gods can fly (Christensen 1990: 649).

³A point about notation: in this chapter I use ' $T = \{H_1, \dots, H_n\}$ ' to mean that T has axioms H_1, \dots, H_n , not that it has theorems H_1, \dots, H_n . And a point about terminology: by ' E confirms H with respect to T ' (or ' E confirms H relative to T ') I will mean that E confirms H when the conjunction

Definition 5.1 Let $T = \{H_1, \dots, H_n\}$. Then evidence E bootstrap confirms T exactly if $T \cup E \neq \perp$ and for each $i \in \{1, \dots, n\}$ the following two conditions hold:

1. there is a $H' \subset T$ such that $H_i \notin H'$ and
 - a. E confirms H_i with respect to H' ; and
 - b. there is possible – but non-actual – evidence E' such that E' disconfirms H_i with respect to H' ;
2. there is no $H'' \subseteq T$ such that E disconfirms H_i with respect to H'' .

I have eight comments on this definition:

(1) The reader should be warned that the definition reflects *my* understanding of Glymour's theory, and that Glymour's book leaves some room for interpretation.⁴ Since my aim is not exegetical I will not argue for the correctness of this interpretation of Glymour's text. Moreover, it is arguable that even if definition 5.1 should fail to capture what Glymour 'really' had in mind, it defines a notion of confirmation that is well worth considering in its own right.

(2) The terms 'confirms' and 'disconfirms' in the clauses of the definition can, as already intimated, be cashed out in more than one way.⁵ For example, they can also be understood in hypothetico-deductive terms, or in probabilistic terms. Thus, this definition is no longer tied to the Hempelian account of confirmation, as is Glymour's original definition.

(3) Coupled with certain confirmation theories, in the face of subclause 1.a the second clause amounts to no more than the requirement that the theory at issue be consistent (and is thus redundant given the requirement that the theory be consistent with the evidence). Suppose for instance that the notions of confirmation and disconfirmation are understood hypothetico-deductively. Without loss of generality, consider a theory consisting of three axioms, $T = \{H_1, H_2, H_3\}$, and suppose that E confirms H_1 with respect to H_2 , and thus that (a) $\{H_1, H_2\} \vdash E$, but disconfirms H_1 with respect to H_3 , and thus that (b) $\{H_1, H_3\} \vdash \neg E$. Then T must be inconsistent. For it follows from (a) that $\{\neg E\} \vdash \neg H_1 \vee \neg H_2$ and, similarly, it follows from (b) that $\{E\} \vdash \neg H_1 \vee \neg H_3$. And thus, by Constructive Dilemma, $\{E \vee \neg E\} \vdash (\neg H_1 \vee \neg H_2) \vee (\neg H_1 \vee \neg H_3)$, or, put differently,

of hypotheses in T is taken as an auxiliary. And similarly, by ' E confirms H with respect to H' ' with $H' \subset T$ I will mean that E confirms H when the conjunction of hypotheses in H' is taken as an auxiliary.

⁴Cf., e.g., Christensen (1997). Pondering various possibilities of how relative confirmation can provide 'real' confirmation, he conjectures that Glymour '[takes] certain complicated structures of interlocking relative confirmation to constitute real confirmation of a set of hypotheses' (*op. cit.*: 372). As may be clear, I think this conjecture is correct. Earman and Salmon (1992: 52 ff) seem to interpret Glymour in the same way as I do.

⁵Or better, the phrases 'confirms with respect to' and 'disconfirms with respect to'; the exact understanding of 'with respect to' will depend on the interpretation of 'confirms'/'disconfirms'.

$\{E \vee \neg E\} \vdash \neg H_1 \vee \neg H_2 \vee \neg H_3$. From which it follows that $\vdash \neg(H_1 \wedge H_2 \wedge H_3)$, i.e., $\vdash \neg T$. However, it will be immediately clear that according to the common probabilistic understanding of confirmation and disconfirmation, if some T and E satisfy clause 1 but not clause 2, that does not entail that T is inconsistent. So if the definition is to be neutral as regards theories of non-bootstrap confirmation, then the second clause cannot be dispensed with.

(4) Glymour originally allowed bootstrap testing in which evidence could confirm a hypothesis relative to itself – what later came to be called ‘macho-bootstrapping’. However, under the pressure of criticism from, among others, Christensen (1983), Edidin (1983) and van Fraassen (1983a) he later restricted the auxiliaries admissible in a test to hypotheses other than the one under scrutiny in that test (the so-called ‘wimp-bootstrapping,’ see Earman and Glymour (1988: 261)), as does clause 1 by requiring that the hypothesis under scrutiny not be in the set of hypotheses from which the auxiliaries in that hypothesis’s test are taken. (Notice that a similar restriction in clause 2 would be superfluous. If a hypothesis is disconfirmed by the evidence, then this seems to be no less – but rather more – damaging to that hypothesis, and hence also to any theory that includes it, when the hypothesis served itself as an auxiliary in that test, than when only other hypotheses did.)

(5) Bootstrap confirmation as presented here is defined for finitely axiomatizable theories only. It is not theoretically impossible to generalize definition 5.1 to the infinite case, but it is hard to see how bootstrap confirmation of such theories could *practically* be achieved. Thus, in what follows, by ‘theory’ I will mean finitely axiomatizable theories.⁶

⁶ Though Glymour does not note this, one must also assume that theories are *naturally* axiomatized (in some sense of ‘natural’) lest the notion of bootstrap confirmation become relative to a given axiomatization. To see this, just consider that since every finitely axiomatizable theory is axiomatizable by just one axiom – given any finite axiomatization, take the conjunction of the axioms – and since, given that I have excluded macho-bootstrapping, a theory with only one axiom cannot be bootstrap-tested, without some notion of a natural axiomatization it may be possible to claim of one and the same theory both that it is and that it is not bootstrap confirmed by the evidence. *With* that notion, we can stipulate that a theory is bootstrap confirmed by the evidence if its natural axiomatization is bootstrap confirmed by the evidence. It seems that the notion of a natural axiomatization has been around in the logical literature for some time. However, only recently an attempt has been made to explicate it; see Gemes (1993) (also his (1994), (1997); Schurz’s (1991) theory of relevant deduction can also be thought of as such an attempt). I do not want to commit myself to Gemes’s or any other explication. For one, according to Gemes’s (1983: 483) definition, a theory will usually be naturally axiomatized by only one axiom, which would not be helpful for the question at hand. Only in a footnote (Gemes 1983: 483n3) an additional condition is suggested, according to which this no longer would be the case.

For present purposes the intuitive notion of a natural axiomatization seems clear enough. Every theory to be presented by its axioms in this chapter, both in the examples and in the proofs of the theorems, is assumed to be naturally axiomatized in this intuitive sense.

(6) As in a bootstrap test of a theory T all the tests of the individual axioms of T rely on auxiliaries that also come from T , it might seem a questionable feature of this account that, provided both clauses of definition 5.1 are satisfied, it allows us to conclude that the evidence confirms the theory, period, and not just that it confirms the theory with respect to itself (a conclusion – note – that would be barely significant unless one is already willing to accept the theory). More than questionable, in fact: a common response of those who first learn about bootstrap testing is to exclaim that the procedure is patently circular. After all – it is said – the very theory the truth of which is at stake in the test is presupposed in that test in the sense that it is allowed to supply the auxiliaries needed in the tests of the separate hypotheses comprised by the theory. At first sight, the situation may indeed seem analogous to one in which we (correctly) derive some proposition A using A itself as a premise and then present that as a proof of A (instead of just as a proof of $A \rightarrow A$, or of $\{A\} \vdash A$). But it is not. Consider: if bootstrap testing were really circular, then how could any theory consistent with the evidence ever fail to be bootstrap confirmed? But if the ‘non-triviality subclause’ 1.b is satisfied, such a theory *can* fail to be bootstrap confirmed. For what the subclause ensures is that adopting certain hypotheses as auxiliaries in testing some other hypothesis does not guard the latter against disconfirmation whatever the data. This, I believe, is as straightforward a way as any to see the non-circularity of bootstrap testing.⁷

(7) Definition 5.1 defines bootstrap confirmation. What about bootstrap disconfirmation? Glymour does not say, and it seems this notion can be defined in more than one plausible way. Like Glymour I will mainly concern myself with bootstrap confirmation. Nevertheless, it will prove useful later on to have a definition of bootstrap disconfirmation at hand. As such I propose this:

Definition 5.2 *Evidence E bootstrap disconfirms theory $T = \{H_1, \dots, H_n\}$ precisely if for at least one $i \in \{1, \dots, n\}$ there is a $H' \subseteq T$ such that E disconfirms H_i with respect to H' .*

(8) My final comment has a heuristic intent. It may be helpful to think of definition 5.1 as indicating some sort of coherence of the axioms of a theory both with one another and with the evidence. In chapter 1, I indicated that intuitively coherence is a matter of how well the various propositions involved hang together and, secondly, that coherence comes in degrees. Clearly, a positive bootstrap test by evidence E of a theory $T = \{H_1, \dots, H_n\}$ is an indication of the hypotheses in

⁷One may insist that the mere fact that a theory is (in a sense) presupposed in its own test is sufficient to make the procedure circular. Of course one may define circularity in any way one likes, but the crucial issue is whether the fact that a theory supplies auxiliaries for testing its own axioms is *vicious*. And I can only challenge anyone who holds that it is to point out why that is so.

T and E hanging together in a very clear sense: the hypotheses help each other to obtain support from the evidence. Of course the definition cannot quite be a definition of coherence, for it does not allow for coherence to be a matter of degree. The quantitative theory of bootstrap confirmation to be developed in this chapter does make graded judgements possible and I will come back to the relationship between the measure of bootstrap confirmation proposed below and our general coherence intuitions in section 5.6.

This concludes my discussion of definition 5.1. As the second point above indicates, the definition is independent of any specific account of non-bootstrap confirmation. Nonetheless, in order to assess the merits of Glymour's theory of bootstrap confirmation, we do need to combine it with such an account. Given the strong criticisms of Glymour's original definition, Hempel's positive instance account does not appear to be a viable option. Moreover, if we accept the claim that a quantitative account of bootstrap confirmation is more preferable than a merely qualitative account, then the most obvious option is to combine definition 5.1 with the Bayesian theory of confirmation.⁸ To that effect we must reinterpret each of the conditions in definition 5.1 within a Bayesian framework. I will consider each in turn:

(1) The first part of the first condition requires that there is a $H' \subset T$ such that $H_i \notin H'$ and E confirms H_i with respect to H' according to some confirmation condition. Naively, it may appear that this condition is satisfied if $p(H|H' \wedge E) > p(H)$. But although in that case the probability of the hypothesis is certainly raised, this may have nothing to do with the evidence since it could also be due to the auxiliary hypotheses H' . Therefore, a better interpretation of this condition is that $p(H|H' \wedge E) > p(H|H')$.

This leads to the following definition:⁹

Definition 5.3 *Evidence E probabilistically bootstrap confirms hypothesis H relative to a complex of auxiliaries H' exactly if $p(H|H' \wedge E) > p(H|H')$.*

For future reference, let me also define the notion of probabilistic bootstrap disconfirmation of a hypothesis:

⁸By choosing a Bayesian approach I simply follow the mainstream in current analytic philosophy. However, it is noteworthy – as an anonymous referee reminded Igor Douven and me – that there exist other quantitative approaches to confirmation besides Bayesianism, such as Shafer's (1976) Dempster-Shafer belief functions, Zadeh's (1978) possibility measures, and Spohn's (1988) ranking functions (see Halpern (2003: Ch. 2) for an excellent overview of the different approaches to represent uncertainty).

⁹As will become clear below, the use of the concept of bootstrap confirmation is different in the case of the evidence probabilistically bootstrap confirming a hypothesis relative to the auxiliaries than in the case of the evidence probabilistically bootstrap confirming a theory. I am confident that this will not lead to confusion.

Definition 5.4 Evidence E probabilistically bootstrap disconfirms hypothesis H relative to a complex of auxiliaries H' exactly if $p(H | H' \wedge E) < p(H | H')$.

(2) The second part of the first condition reads that there should be possible – but non-actual – evidence E' such that E' disconfirms H_i with respect to H' . This would translate as the requirement that there is an E' such that $p(H | H' \wedge E') < p(H | H')$. However, this need not be listed separately for it follows directly from definitions 5.3 and 5.4. For if $p(H | H' \wedge E) > p(H | H')$, then there is possible evidence E' such that $p(H | H' \wedge E') < p(H | H')$, that is, if E bootstrap confirms H_i relative to H' , then there is possible evidence disconfirming H relative to the same auxiliary or auxiliaries as that or those relative to which E confirms it. After all, by the law of total probability, we have that

$$p(H | H') = p(H | H' \wedge E)p(E | H') + p(H | H' \wedge \neg E)p(\neg E | H'). \quad (5.1)$$

Thus, since $p(E | H') = 1 - p(\neg E | H')$, the probability $p(H | H')$ is a mixture of $p(H | H' \wedge E)$ and $p(H | H' \wedge \neg E)$.¹⁰ And so if $p(H | H' \wedge E) > p(H | H')$, it must be that $p(H | H' \wedge \neg E) < p(H | H')$.

(3) The last condition of definition 5.1 was that there is no $H'' \subseteq T$ such that E disconfirms H_i with respect to H'' . Using definition 5.4, this condition translates into the condition that there is no $H'' \subseteq T$ such that $p(H_i | H'' \wedge E) < p(H_i | H'')$.

Summing up these conditions yields the following definition for probabilistic bootstrap confirmation:^{11,12}

Definition 5.5 Evidence E probabilistically bootstrap confirms a theory $T = \{H_1, \dots, H_n\}$ precisely if $p(T \wedge E) > 0$ and for each $i \in \{1, \dots, n\}$ it holds that

1. there is a $H' \subset T$ such that $H_i \notin H'$ and $p(H_i | H' \wedge E) > p(H_i | H')$; and

¹⁰Here it may be helpful to note that if $p(H | H' \wedge E) > p(H | H')$, it must hold that both $0 < p(H' \wedge E) < 1$ and $p(H' \wedge E) \neq p(H')$, and hence also that $0 < p(E | H') < 1$.

¹¹ According to Duhem, Quine, and others, confirmation *generally* is three-place. Since I know of no air-tight argument showing that confirmation is *necessarily* three-place, it seems best to at least formally leave open the possibility that evidence confirms a hypothesis relative to the empty set, i.e., without the aid of any auxiliaries. Strictly speaking, the clauses of definition 5.5 make no sense in case the subsets of T they refer to are empty. It should be obvious, however, that in that case ' $p(H_i | H' \wedge E) > p(H_i | H')$ ' is to be read as $p(H_i | E) > p(H_i)$; similarly for ' $p(H_i | H'' \wedge E) < p(H_i | H'')$ ' in the second clause.

¹²As an anonymous referee brought to the attention of Igor Douven and me, it is insufficient to require that $T \cup \{E\} \neq \perp$ (as is done in definition 5.1, and as we did in definition 5.5 in an earlier version of the (2005b) paper) given that T may be consistent with E and yet it may hold that $p(T \wedge E) = 0$ (unless one assumes all probability functions to be strict, which I don't); and of course one would not want to say that a theory can be confirmed in any sense by evidence conditional on which it has probability 0.

2. there is no $H'' \subseteq T$ such that $p(H_i | H'' \wedge E) < p(H_i | H'')$.

For probabilistic bootstrap disconfirmation, I propose the following definition:¹³

Definition 5.6 *Evidence E probabilistically bootstrap disconfirms theory $T = \{H_1, \dots, H_n\}$ iff for at least one $i \in \{1, \dots, n\}$ there is a $H' \subseteq T$ such that $p(H_i | H' \wedge E) < p(H_i | H')$.*

This concludes my discussion of Glymour's theory of bootstrap confirmation. The next section will propose a measure of confirmation that is based on the definition of probabilistic bootstrap confirmation.

Unless stated otherwise, from now on I will understand 'bootstrap confirmation' and 'bootstrap disconfirmation' as 'probabilistic bootstrap confirmation' and 'probabilistic bootstrap disconfirmation.'

5.3 A Class of Measures of Bootstrap Confirmation

Before I can begin to formulate a measure of bootstrap confirmation, I should start by considering what the desiderata for such a measure are. That is, I must make clear which relations between the hypotheses in the theory and the evidential statements matter intuitively for the measure in which the evidence bootstrap confirms a theory. I suggest that the following four desiderata are quite sensible (although there may be others):

Firstly, if the hypotheses a theory T consists of are all bootstrap confirmed (according to definition 5.3) to a higher degree than those belonging to another theory T' , then, all else being equal, the first theory should have a higher degree of bootstrap confirmation.

The second desideratum is inspired by an argument Glymour (1980a: 139–142) presents against Bayesianism. According to him, Bayesianism cannot express the value of variety of evidence. For Glymour (1980a: 140) '[w]hat makes one way of testing relevantly different from another is that the hypotheses used in one computation are different from the hypotheses used in the other computation.' I am not sure whether this is the best description of variety of evidence, let alone that it is the only possible one, but as a general desideratum it seems quite sensible. As a second desideratum, therefore, I would propose that given two theories, T and T' , if T is tested with respect to a greater variety of (complexes

¹³Again, I only give this definition for later purposes. Specifically, I am not arguing that it mirrors best our intuitions with respect to bootstrap disconfirmation. I simply need this definition in the construction of some of the theorems concerning the measure of bootstrap confirmation to be proposed later in this chapter. For heuristical purposes, I have labeled it probabilistic bootstrap disconfirmation.

of) auxiliaries than T' , then, all else being equal, E should bootstrap confirm T to a greater extent than it bootstrap confirms T' .

Thirdly, a measure of bootstrap confirmation should increase if a larger number of propositions are bootstrap confirmed by the evidence and some auxiliaries.

Lastly, it seems that the measure should allow for degrees in which hypotheses are shielded from disconfirmation by other hypotheses. For while this shielding from disconfirmation in Glymour's approach can only be of the yes-or-no type, this is no longer so in a probabilistic setting. Instead, it seems that a measure of bootstrap confirmation should reflect the degree to which hypotheses are shielded by the auxiliaries.

It can be readily appreciated that we can satisfy the desiderata by summing over the degree of confirmation all hypotheses receive from the evidence relative to all possible combinations of available auxiliaries. Therefore, the measure we are after has the form of a double summation, first over all the possible sets of auxiliaries with respect to which a hypothesis may be confirmed and then over the total bootstrap confirmation each of the hypotheses receives from all sets of auxiliaries and the evidence.

Evidently, for this to be feasible, we need a measure for the degree to which a hypothesis H_i is confirmed by any of the subsets of $T \setminus \{H_i\}$ and the evidence. Therefore, we need to adapt the Bayesian measures of confirmation so as to make them capable of measuring the confirmation of a hypothesis by the evidence relative to some auxiliaries. In this chapter I will limit my attention to the following four measures of confirmation:

- the difference measure d : $d(H, E) =_{df} p(H|E) - p(H)$;
- the log-ratio measure r : $r(H, E) =_{df} \log [p(H|E)/p(H)]$;
- Carnap's relevance measure c : $c(H, E) =_{df} p(H \wedge E) - p(H)p(E)$.
- the log-likelihood measure l : $l(H, E) =_{df} \log [p(E|H)/p(E|\neg H)]$;

Interestingly, it seems that not all of these measures can be satisfactorily adapted to measuring bootstrap confirmation. For if a measure is to measure the degree to which evidence E bootstrap confirms H relative to a complex of auxiliaries H' , it seems such a measure should equal 0 (or at least have a neutral value) if the hypothesis is independent of the evidence and the auxiliaries and that it should be larger than zero (or larger than the neutral value) if the hypothesis is bootstrap confirmed by the evidence relative to the auxiliaries and negative (smaller than the neutral value) if the hypothesis is bootstrap disconfirmed by the evidence relative to the auxiliaries. By definition 5.3, a hypothesis is bootstrap confirmed by the evidence relative to some auxiliary hypotheses if and only if $p(H|H' \wedge E) > p(H|H')$. The difference measure, the ratio measure and Carnap's relevance measure can all be readily adapted to satisfy this demand. Here are three proposals to that effect:

- $\delta(H; H', E) =_{df} p(H|H' \wedge E) - p(H|H')$;
- $r(H; H', E) =_{df} \log [p(H|H' \wedge E)/p(H|H')]$;
- $c(H; H', E) =_{df} p(H \wedge H' \wedge E) - p(H|H')p(H' \wedge E)$.

The following theorem shows that each of these three measures satisfies the above requirement (for a proof, see Appendix 5 A):

Theorem 5.1 *For any hypothesis H , evidence E and complex of auxiliaries H' , it holds that $\delta(H; H', E)$, $r(H; H', E)$ and $c(H; H', E)$ are zero when H is independent of E and H' and positive (negative) if H is bootstrap confirmed (disconfirmed) by E with respect to H' .*

But this is not the case for the most plausible adaptation of the log-likelihood measure:

$$- \mathfrak{l}(H; H', E) =_{df} \log [p(H' \wedge E|H)/p(H' \wedge E|\neg H)];$$

It can easily be checked that if a hypothesis is independent of the auxiliaries and the evidence, $\mathfrak{l}(H; H', E) = 0$. But to satisfy the above requirement, it should also be the case that if H is bootstrap confirmed by E relative to H' , the value of $\mathfrak{l}(H; H', E)$ is positive. However, as the following theorem shows, this is not always the case (for a proof, see Appendix 5 A):

Theorem 5.2 *Given a hypothesis H , evidence E and complex of auxiliaries H' it is not always the case that $\mathfrak{l}(H; H', E) > 0$ if H is bootstrap confirmed by E relative to H' .*

On account of this result, I will focus on the other three measures adapted to bootstrapping, which I will call the bootstrap difference measure, the bootstrap ratio measure and the bootstrap relevance measure. With the help of these it is relatively straightforward to define a measure of bootstrap confirmation. Let ' $\{T_i\}$ ' denote the class of all finitely axiomatizable, deductively closed theories that can be formulated in a given language and ' $\{E_j\}$ ' the class of sentences of that language apt to report evidence (which I will assume to be coextensive with the class of all sentences). This leads to an expression of the measure I am after as a function $\mathcal{B}(T, E): \{T_i\} \times \{E_j\} \rightarrow \mathbb{R}$. Next, let ' \mathfrak{m} ' denote one of the measures of confirmation adapted to bootstrapping which can then be expressed as the function $\mathfrak{m}(H; H', E): \{H_i\} \times \{E_j\} \rightarrow \mathbb{R}$, with ' $\{H_i\}$ ' the class of all hypotheses.

Now let $T = \{H_1, \dots, H_n\}$. Disallowing macho-bootstrapping, there are for each $H_i \in T$ exactly 2^{n-1} sets of auxiliary hypotheses also in T with respect to which it can be tested, namely, all the elements of the power set of T minus H_i , that is, $\wp(T \setminus \{H_i\})$. Given some ordering $\langle H_{i_1}^T, \dots, H_{i_{2^{n-1}}}^T \rangle$ of $\wp(T \setminus \{H_i\})$, let ' $H_{i_j}^T$ '

denote the j th member of that ordering. Then the following defines a family of three measures of bootstrap confirmation:

Definition 5.7 *If a theory $T = \{H_1, \dots, H_n\}$ is bootstrap confirmed by evidence E according to definition 5.5, then the measure in which E bootstrap confirms T is given by:*

$$\mathcal{B}_m(T, E) =_{df} \sum_{i=1}^n \sum_{j=1}^{2^{n-1}} m(H_i; H_{ij}^T, E), \quad (5.2)$$

for $m \in \{\delta, r, c\}$.

I have three remarks on the definition of the measure of bootstrap confirmation:

(1) I have explicitly demanded that T be bootstrap confirmed to avoid the rather awkward consequence that a theory can be bootstrap disconfirmed in the qualitative sense of definition 5.6, while being bootstrap confirmed to a positive degree (that this is possible is shown by theorem 5.5 below). Alternatively, I could also simply resist making the connection between a positive value for any of the \mathcal{B}_m and the evidence bootstrap confirming a theory. However, it will appear below that in that case the problem of independent hypotheses to be discussed in section 5.5 becomes much more serious.

(2) Although I have given a definition of bootstrap disconfirmation, I will not give a measure of it. This is primarily because it is not directly clear what such a measure would look like and how it could be distinguished from a measure of bootstrap confirmation. Explicating bootstrap disconfirmation would be one of the projects this chapter invites.

(3) I have not included any weighing factors in definition 5.7. One reason for this is that from the desiderata above it follows that \mathcal{B}_m should not be defined as a straight average of the different values of $m \in \{\delta, r, c\}$, for in that case the second desideratum would no longer be satisfied since it would no longer hold generally that if a hypothesis is bootstrap confirmed with respect to a larger number of auxiliaries, then the total value of bootstrap confirmation increases (for one way in which this may arise is by adding additional axioms to a theory, and it is easy to see that in that case the degree of bootstrap confirmation defined as the straight average of the different values of $m \in \{\delta, r, c\}$ may decrease). Nevertheless, it may turn out to be necessary to add some other weighing factors later on.

To get a feel for definition 5.7, it may be helpful to see an actual application of it. Consider the following example in which the bootstrap support the evidence supplies to a theory consisting of four hypotheses is calculated:

Example 5.1 Theory T consists of axioms H_1, H_2, H_3 , and H_4 . Each of these hypotheses has a prior probability of .25, and they are mutually probabilistically independent. Evidence E has a prior probability of .5. Further, the following conditions hold:¹⁴

- $p(E \wedge H_i) = .125$, for all $i \in \{1, \dots, 4\}$;
- $p(E \wedge H_i \wedge H_j) = .05$, for all $i, j \in \{1, \dots, 4\}$ such that $i \neq j$;
- $p(E \wedge H_i \wedge H_j \wedge H_k) = .015$, for all $i, j, k \in \{1, \dots, 4\}$ such that $i \neq j \neq k$;
- $p(E \wedge H_1 \wedge \dots \wedge H_4) = .0038$.

To calculate the measure of support E provides to T , first choose a specific measure one will use. I will use the bootstrap difference measure δ . In that case, for all $i \in \{1, \dots, 4\}$ it holds that

- $\delta(H_i, E) = \frac{p(H_i \wedge E)}{p(E)} - p(H_i) = .125/.5 - .25 = 0$;

- for all $j \in \{1, \dots, 4\}$ such that $i \neq j$:

$$\delta(H_i; H_j, E) = \frac{p(H_i \wedge H_j \wedge E)}{p(H_j \wedge E)} - \frac{p(H_i \wedge H_j)}{p(H_j)} = \frac{.05}{.125} - \frac{.0625}{.25} = .15;$$

- for all $j, k \in \{1, \dots, 4\}$ such that $i \neq j \neq k$:

$$\begin{aligned} \delta(H_i; H_j \wedge H_k, E) &= \frac{p(H_i \wedge H_j \wedge H_k \wedge E)}{p(H_j \wedge H_k \wedge E)} - \frac{p(H_i \wedge H_j \wedge H_k)}{p(H_j \wedge H_k)} \\ &= \frac{.015}{.05} - \frac{.015625}{.0625} = .05; \end{aligned}$$

- for all $j, k, l \in \{1, \dots, 4\}$ such that $i \neq j \neq k \neq l$:

$$\begin{aligned} \delta(H_i; H_j \wedge H_k \wedge H_l, E) &= \frac{p(H_i \wedge H_j \wedge H_k \wedge H_l \wedge E)}{p(H_j \wedge H_k \wedge H_l \wedge E)} \\ &\quad - \frac{p(H_i \wedge H_j \wedge H_k \wedge H_l)}{p(H_j \wedge H_k \wedge H_l)} \\ &= \frac{.0038}{.015} - \frac{.00390625}{.015625} \approx .003. \end{aligned}$$

¹⁴Here and in the rest of this chapter, probability functions will be specified without a proof that they are probability functions. It is nowadays easy to check that they are, however, by means of the function **InequalityInstance** of *Mathematica* (versions 4.1 and higher); see Fitelson (2001a: 93–100) for an explanation of how to do this. Alternatively, one may use the PrSAT package for *Mathematica* (version 5), which can be downloaded from <http://www.fitelson.org>.

Since for each H_i there is exactly one way in which it can be tested relative to no auxiliary hypotheses, three different ways in which it can be tested relative to one auxiliary hypothesis, three different ways in which it can be tested relative to two auxiliary hypotheses, and one way in which it can be tested relative to three auxiliary hypotheses, the total bootstrap support each of the H_i individually receives from E equals (approximately): $(1)(0) + (3)(.15) + (3)(.05) + (1)(.003) = .603$. Since the bootstrap support T receives is just the sum of the bootstrap support each of its axioms receives, $\mathcal{B}_\delta(T, E) \approx (4)(.603) = 2.412$, with \mathcal{B}_δ the value of $\mathcal{B}_m(T, E)$, when calculated by using the bootstrap difference measure. ♦

As a conclusion to this section, let me present a number of theorems concerning the relation between the measures of coherence \mathcal{B}_m ($m \in \{\delta, r, c\}$) and definitions 5.5 and 5.6 (for the proofs, see Appendix 5 A).

Theorem 5.3 *For all T and E , if E bootstrap confirms T , then $\mathcal{B}_m(T, E) > 0$ for all $m \in \{\delta, r, c\}$.*

Theorem 5.4 *For all T and E and for all $m \in \{\delta, r, c\}$, if $\mathcal{B}_m(T, E) < 0$, then E bootstrap disconfirms T .*

The next two theorems show that the converses of the above two theorems do not hold.

Theorem 5.5 *There is no $a \in \mathbb{R}$ and no $m \in \{\delta, r, c\}$ such that, for all T and E , if $\mathcal{B}_m(T, E) > a$, then E bootstrap confirms T .*

Theorem 5.6 *There is no $a \in \mathbb{R}$ and no $m \in \{\delta, r, c\}$ such that, for all T and E , if E bootstrap disconfirms T , then $\mathcal{B}_m(T, E) < a$.*

Finally, the following theorem considers the relation between two different theories:

Theorem 5.7 *For all T, T' , and E , if $T \subset T'$ and E does not bootstrap disconfirm T' , then $\mathcal{B}_m(T', E) \geq \mathcal{B}_m(T, E)$ for all $m \in \{\delta, r, c\}$; if in addition E bootstrap confirms T' , then $\mathcal{B}_m(T', E) > \mathcal{B}_m(T, E)$ for all $m \in \{\delta, r, c\}$.*

Each of the above theorems applies to all members of the family of measures of bootstrap confirmation. Nonetheless, it should not be concluded from this that the measures of bootstrap confirmation do not differ significantly. As the next section will show, there are some important differences between the three measures, although it will remain unclear which is the most satisfactory one.

5.4 Comparing the Different Measures

The similarities between the measures of bootstrap confirmation defined above and those of coherence as mutual support presented in chapter 2 are obvious. Both are measures of mutual support, both consider a number of confirmation relations between different parts of a theory and/or the evidence, and both define a family of different measures, each of which is based on an adaptation of a popular Bayesian measure of confirmation. In chapter 2, I argued that we have good intuitive grounds for preferring measure C_d over the other two measures. In this section, I will not make a similar claim. Instead, I will discuss a few important differences between the three measures of bootstrap confirmation and leave the answer to the question of which measure has the most satisfactory consequences to the reader. The reason for this is that I am not sure what conclusions to draw from the differences discussed in this section.

5.4.1 Threshold Values

Note that theorem 5.5 only indicates that there is no *general* numerical threshold value for qualitative bootstrap confirmation in the sense that, for any T and E , if one knows that $\mathcal{B}_m(T, E)$ for a given $m \in \{\delta, r, c\}$ has a value above that threshold, one can immediately infer that E bootstrap confirms T . This leaves open the possibility that with every *particular* theory T some value a can be associated such that, if $\mathcal{B}_m(T, E) > a$ for some E and $m \in \{\delta, r, c\}$, then E bootstrap confirms T . The following theorems will show that this possibility holds for both the bootstrap difference measure and the bootstrap relevance measure, but not for the bootstrap ratio measure. For the first two, something stronger even holds. For in those cases there is a threshold that depends only on the number of axioms in the theory (for a proof, see Appendix 5 B):

Theorem 5.8 *For all $n \in \mathbb{N}$, T , and E , if $T = \{H_1, \dots, H_n\}$ and $\mathcal{B}_m(T, E) \geq n 2^{n-1} - 1$ for $m \in \{\delta, c\}$, then E bootstrap confirms T .*

However, the same does not hold for the bootstrap ratio measure (for a proof, see Appendix 5 B):

Theorem 5.9 For all $n \in \mathbb{N}$, T , and E , if $T = \{H_1, \dots, H_n\}$, $\mathcal{B}_r(T, E)$ can be arbitrarily large while E still bootstrap disconfirms T .

5.4.2 A Second Difference Between the Three Bootstrap Measures of Confirmation

The last subsection only showed that \mathcal{B}_δ and \mathcal{B}_c differ from \mathcal{B}_r , not that \mathcal{B}_δ and \mathcal{B}_c differ from each other. However, \mathcal{B}_δ and \mathcal{B}_c , too, have some strikingly different consequences. For consider the following example:

Example 5.2 Theory T has as axioms hypotheses H_1 , H_2 , and H_3 , each of which has a prior probability of .1; E has a prior probability of .5. Also the following hold:

- $p(H_i \wedge H_j) = .015$ for all $i, j \in \{1, 2, 3\}$ such that $i \neq j$;
- $p(H_1 \wedge H_2 \wedge H_3) = .005$;
- $p(E \wedge H_i) = .05$ for all $i \in \{1, 2, 3\}$;
- $p(E \wedge H_i \wedge H_j) = x$ for all $i, j \in \{1, 2, 3\}$ such that $i \neq j$ and x a variable;
- $p(E \wedge H_1 \wedge H_2 \wedge H_3) = a$ with a a positive constant;
- for no H_i ($i \in \{1, 2, 3\}$) does E bootstrap disconfirm H_i relative to any subset of T ; therefore, and because $x \leq p(H_i \wedge H_j)$ ($i, j \in \{1, 2, 3\}, i \neq j$) and $a \leq p(H_1 \wedge H_2 \wedge H_3)$, it holds that $.0075 \leq x \leq .015$ and $.333x \leq a \leq .005$.¹⁵

Using the bootstrap difference measure yields the following expression for $\mathcal{B}_\delta(T, E)$:

$$\begin{aligned} \mathcal{B}_\delta(T, E) &= 3 \left[2 \left(\frac{x}{.05} - \frac{.015}{.1} \right) + \frac{a}{x} - \frac{.005}{.015} \right] \\ &= 120x + \frac{3a}{x} - 1.9. \end{aligned}$$

Evaluation of this function shows that it reaches its minimum at

$$x_0 = \sqrt{\frac{a}{40}}. \quad (5.3)$$

For this model, it turns out that one can choose values for a such that $x_0 > .0075$ (the lower bound on x). Given such values, and given that x_0 is a minimum, $\mathcal{B}_\delta(T, E)$ will first decrease until it reaches x_0 , and only then increase – see, e.g., the graph in figure 5.1, which gives the value of $\mathcal{B}_\delta(T, E)$ as a function of x for $a = .005$. This means that, all else being equal, raising the probability of a hypothesis conditional on the evidence and some $T' \subset T$ will sometimes have the effect of lowering the value of $\mathcal{B}_\delta(T, E)$. ♦

¹⁵If $x < .0075$, then $\delta(H_i, H_j \wedge E) < 0$ (for all $i, j \in \{1, 2, 3\}, i \neq j$). Likewise, if $a < .333x$, then $\delta(H_i; H_j \wedge H_k, E) < 0$ (for all $i, j, k \in \{1, 2, 3\}, i \neq j \neq k$).




Figure 5.1: Value of $\mathcal{B}_S(T, E)$ as a function of $x : .0075 \leq x \leq .015$

As the following two examples show, the above result does not hold for either the bootstrap ratio measure or the bootstrap relevance measure.

Example 5.3 *Take everything equal to example (5.2). But now measure the degree of bootstrap confirmation by means of $\mathcal{B}_r(T, E)$ instead of $\mathcal{B}_S(T, E)$. Expressed as a function of x , the former becomes:*

$$\mathcal{B}_r(T, E) = 3 \left(2 \log \left[\frac{.1}{.015} \frac{x}{.05} \right] + \log \left[\frac{.015}{.005} \frac{a}{x} \right] \right) \quad (5.4)$$

$$= 3 \left(2 \log \left[\frac{400x}{3} \right] + \log \left[\frac{3a}{x} \right] \right). \quad (5.5)$$

The derivative of this expression is given by:

$$\frac{d\mathcal{B}_r(T, E)}{dx} = \frac{3}{x}. \quad (5.6)$$

For all $x > 0$, this derivative is positive and, therefore, $\mathcal{B}_r(T, E)$ increases as x increases. ♦

Example 5.4 *Take everything equal to example (5.2). But now use $\mathcal{B}_c(T, E)$ to measure the degree of bootstrap confirmation, which is given by:*

$$\begin{aligned} \mathcal{B}_c(T, E) &= 3(2x - .015) + 3(a - .333x) \\ &= 5x + 3a - .045. \end{aligned}$$

It is plain that, for any fixed value of a , this function increases linearly with x .

This result can even be generalized. For it is easily seen that for any $n \in \mathbb{N}$, if $T = \{H_1, \dots, H_n\}$, then for any $m: 1 \leq m \leq n$, if $H_{i_j} \in T$ for all $j: 1 \leq j \leq m$, $\mathcal{B}_c(T, E)$ is a linearly increasing function of $p(E \wedge H_{i_1} \wedge \dots \wedge H_{i_m})$, all else being equal. ♦

5.4.3 Choosing Between the Measures

Clearly, the three measures of bootstrap confirmation have some notably different consequences. Nonetheless, I do not feel that either of the two consequences provides a definite point in favor of one (or more) of these measures. Firstly, with respect to the first point, it seems that it does provide a point in favor of a measure of bootstrap confirmation m^* if there is a threshold value for B_{m^*} above which one can be sure that a theory is bootstrap confirmed. However, the threshold value for $B_{\delta}(T, E) \geq n 2^{n-1} - 1$ seems so high that its usefulness is limited.

With respect to the example in which the value of $p(E \wedge H_i \wedge H_j)$ was increased while $p(E \wedge H_1 \wedge H_2 \wedge H_3)$ and $p(H_i \wedge H_j)$ were kept constant, it is unclear to me whether any of the B_m ($m \in \{\delta, r, c\}$) should be a monotonically increasing function of $p(E \wedge H_i \wedge H_j)$. For one, it is an elementary consequence of probability theory that

$$p(E \wedge H_i \wedge H_j) = p(E \wedge H_i \wedge H_j \wedge H_k) + p(E \wedge H_i \wedge H_j \wedge \neg H_k), \quad (5.7)$$

and therefore, keeping $p(E \wedge H_1 \wedge H_2 \wedge H_3)$ constant while increasing $p(E \wedge H_i \wedge H_j)$ will have the effect of increasing $p(E \wedge H_1 \wedge H_2 \wedge \neg H_3)$, which – one may feel – should cause a reduction in the degree of bootstrap confirmation. This would not directly explain the minimum in the graph of B_{δ} , but it may indicate that one should not necessarily expect the graph to increase monotonically.

Therefore, pending any intuitively clearer examples, it seems that there is no reason for preferring one of the three measures of bootstrap confirmation over the others.

5.5 Anticipated Objections

In this section I will consider a number of objections that can be raised against the theory of probabilistic bootstrap confirmation in general and the measures of the degree of probabilistic bootstrap confirmation in particular.

5.5.1 Edidin's and Horwich's Arguments Against Bootstrapping

To many it may appear strange that nobody has yet attempted to provide a Bayesian analysis of Glymour's theory of bootstrap confirmation, especially because Glymour (1980b) explicitly lists it as one of the projects that his theory of bootstrap confirmation invites. Part of the reason for this must surely be that Glymour's (1980a) presented a number of important problems for Bayesianism – like, for instance, the problem of old evidence (see Glymour 1980a: 85–93), the

solving of which to many Bayesians may have seemed much more pressing than explicating probabilistic bootstrap confirmation. But another plausible reason is that many Bayesian philosophers seemed to believe that bootstrapping did not need a special Bayesian analysis since it follows directly from Bayesian confirmation theory. For instance, Roger Rosenkrantz (1983: 69) has argued that Glymour ‘is a Bayesian [and] more so than many who march under that banner.’ And Aron Edidin has argued that after a number of necessary corrections to Glymour’s definition of bootstrap confirmation (which need not concern us here)

[w]hat was advanced as an alternative to the Bayesian theory turns out to be derivable from the theory. The principal advantage Glymour claims for his strategy is that it explains how evidence selectively confirms some hypotheses but not others. He argues that Bayesian theorists cannot do the same without recourse to ad hoc restrictions on prior probabilities. But if I’m right about the need to modify the strategy as I suggest, and about its subsequent derivability from the Bayesian theory, then it may be that the principal benefit of the strategy is that it shows how Bayesians can explain the selective relevance of evidence without ad hoc restrictions. (Edidin 1983: 53–54)

If the arguments by Rosenkrantz and Edidin are correct, it may be the case that my theory of bootstrap confirmation, too, can be derived from standard Bayesian confirmation theory. To evaluate this question, let evidence E probabilistically confirm theory T if $p(T | E) > p(T)$ and let E probabilistically disconfirm T if $p(T | E) < p(T)$. The next two theorems show that there is at least some relation between probabilistic confirmation and bootstrap confirmation (the proofs are given in Appendix 5 C).

Theorem 5.10 *For all T and E , if E bootstrap confirms T , then E also probabilistically confirms T .*

Theorem 5.11 *For all T and E , if E probabilistically disconfirms T , then E also bootstrap disconfirms T .*

That is, if a theory is bootstrap confirmed by the evidence it is also probabilistically confirmed by it and if it is probabilistically disconfirmed by the evidence it is also bootstrap disconfirmed by it. However, it cannot be the case that my measure of bootstrap confirmation is derivable from a Bayesian measure of support, for the next two theorems show that neither theorem 5.10 nor theorem 5.11 can be strengthened to a bi-implication (the proofs are given in Appendix 5 C).

Theorem 5.12 *It is not the case that, for all T and E , if E probabilistically confirms T , then E also bootstrap confirms T .*

Theorem 5.13 *It is not the case that, for all T and E , if E bootstrap disconfirms T , then E also probabilistically disconfirms T .*

Alternatively, it may also be the case that my measure of bootstrap confirmation coincides with one of the above measures of confirmation or indeed with any of the other known measures of confirmation (see note 10 in chapter 2). Also, it may be that the degree of bootstrap support provided by some piece of evidence to a theory is just the degree of confirmation of the theory by the evidence given some of those measures of confirmation modulo some scale transformation. This is not the case, however. Note that each of the known measures of confirmation makes the degree of confirmation a theory T receives from evidence E a function of some subset of $\{p(T), p(E), p(T | E), p(T | \neg E), p(E | T), p(E | \neg T)\}$. Call any measure, whether or not actually proposed, that is a function of any such subset a non-bootstrap measure of confirmation. Then we have the following theorem (see Appendix 5 C for a proof):

Theorem 5.14 *For all $m \in \{\delta, r, c\}$ there is no function f such that, for all T and E , $B_m(T, E) = f \circ m(T, E)$, with m any non-bootstrap measure of confirmation.*

Therefore, there is no simple relation between bootstrap confirmation (qualitative or quantitative) and Bayesian confirmation theory.

5.5.2 The Problem of Independent Hypotheses

It may seem that generally adding a hypothesis to a theory that is independent of that theory should not increase its degree of confirmation. However, adding an independent hypothesis can increase a theory's degree of bootstrap confirmation. For consider the following example:

Example 5.5 *Take everything equal to example (5.1). But now add a fifth hypothesis, H_5 , such that it is independent of theory $T = \{H_1, H_2, H_3, H_4\}$ and such that, furthermore, it is not bootstrap confirmed by the evidence relative to any non-empty subset of T . Suppose also that the evidence bootstrap confirms H_5 relative to the empty set. In that case $p(H_5 | H' \wedge E) = p(H_5 | E)$ for all $H' \subset T$. Also, $p(H_5 | H') = p(H_5)$ for all $H' \subset T$. Now assume furthermore that $p(H_5) = .25$ and that $p(H_5 \wedge E) = .15$. In that case, the level of bootstrap support for the new theory $T' = \{H_1, \dots, H_5\}$ as measured with the bootstrap difference measure is approximately $B_\delta(T, E) \approx 5.624$ – much higher than it was in the original case. ♦*

The result can even be generalized, for from the proof of theorem 5.5 it follows directly that adding an independent hypothesis to a theory will *always* increase the value of \mathcal{B}_m for all $m \in \{\delta, r, c\}$ and therefore, if a hypothesis is confirmed by the evidence and if a theory is bootstrap confirmed by the evidence, adding that hypothesis to the theory will always increase the theory's degree of bootstrap confirmation. To many this may seem quite counterintuitive: if a hypothesis is independent of a theory, then how could it increase its degree of confirmation (bootstrap or otherwise)? Nonetheless, there are a few possible replies.

Firstly, the hypothesis is still confirmed by the evidence. If this condition is not satisfied, then a theory's degree of bootstrap confirmation cannot increase, since the theory is no longer bootstrap confirmed according to definition 5.1. But since in the above example H_5 is confirmed by the evidence, it is no longer clear that the bootstrap confirmation of $T' = \{H_1, H_2, H_3, H_4, H_5\}$ should not be larger than that of $T = \{H_1, H_2, H_3, H_4\}$. Instead, it seems to be a natural consequence of the desideratum that the degree of bootstrap confirmation should increase if a larger number of propositions are bootstrap confirmed by the evidence and some auxiliaries (the third desideratum; see section 5.3).

Also, if we take seriously the remark above (in section 5.2) that bootstrap confirmation can be seen as indicating some sort of coherence between the axioms of a theory and the evidence, then it seems quite natural that the degree of bootstrap confirmation should increase. For clearly, hypothesis H_5 hangs together with the evidence – it is confirmed by it – and at least does not undermine any of the other axioms of the theory. Therefore, there seems to be little reason why the coherence between the two should not increase.

5.5.3 Van Fraassen's Critique of Bootstrap Confirmation

What is the connection between bootstrap support and justified belief? Let me consider a puzzle raised in van Fraassen (1983b) which seems to show that qualitative bootstrap confirmation of some theory does not give reason to believe that theory. As will be seen, the quantitative theory of bootstrap confirmation presented above may give rise to basically the same puzzle.

Consider a simple, comparative definition of justification in terms of bootstrap confirmation:

- (B1) If E bootstrap confirms T but not T' , then, if E is our total evidence, belief in T is more justified than belief in T' .
- (B2) If E bootstrap confirms T to a higher degree than T' , then, if E is our total evidence, belief in T is more justified than belief in T' .

Though appealingly simple, this appears to conflict with the following principle:

- (P) If $p(T|E) > p(T'|E)$, then, if E is our total evidence, belief in T' cannot be more justified than belief in T ,

a principle that van Fraassen (1983b) presents as a truism.

Van Fraassen's argument that (B1) and (P) conflict goes roughly as follows: suppose $T = \{H_1, H_2, H_3\}$ is bootstrap confirmed by E (according to definition 5.5). Suppose in particular that E confirms H_1 relative to H_2 (and perhaps also relative to H_3 and to $H_2 \wedge H_3$), E confirms H_2 relative to H_3 (and perhaps also relative to $H_1 \wedge H_3$, but not relative to H_1), and E confirms H_3 relative to H_1 (and perhaps also relative to $H_1 \wedge H_2$). Now the subset $T' \subset T$, containing only H_1 and H_2 , is *not* bootstrap confirmed by E (for E does not confirm H_2 relative to H_1). Without loss of generality, assume that $\{H_1, H_2, E\} \not\models H_3$. Then it is an elementary truth of probability theory that $p(T'|E) > p(T|E)$. So, according to (B1) we are more justified in believing T than in believing T' , but according to (P) we are not.

It must be immediately clear that although (B1) always conflicts with (P), this does not hold for (B2). In quantitative bootstrapping there are various ways in which a theory can become more bootstrap confirmed. Many of these will simply raise the probability of the theory given the evidence, and in such cases (B2) and (P) are in perfect accordance with each other. This said, it is not hard to see that in other cases (B2) and (P) will still conflict.

For example, consider a subset T' of T of example 5.1, consisting of hypotheses H_1, H_2, H_3 . This theory has a probability of .03 given evidence E , which clearly exceeds that of T given E (= .0076). However, it is bootstrap-supported by E to a much lower degree than T , namely $B(T', E) = 1.05$ (while $B(T, E) \approx 2.412$). Thus, according to (B2) we are more justified in believing T than in believing T' , but according to (P) we are not. Hence, degree of bootstrap confirmation and probability may pull in opposite directions; they cannot always be jointly maximized. But then how can bootstrap confirmation be related to justification?

Van Fraassen's conclusion is that it cannot. This is not to say that van Fraassen believes a positive bootstrap test is insignificant. Quite the contrary – he believes it gives reason to *accept* a theory, where the notion of acceptance is considerably weaker than that of belief (van Fraassen 1983a and 1983b). More specifically, acceptance of a theory involves the belief that it is empirically adequate (roughly, true of the observable part of the world) as well as a commitment to use the theory's conceptual apparatus in describing future phenomena (van Fraassen 1980: Ch. 1). Should his view on bootstrap testing be correct, then that hardly detracts from the importance of having a quantitative account of bootstrap confirmation: surely it makes sense to say that one bootstrap test provides stronger reason to accept a given theory than another bootstrap test, and it seems that only a quantitative bootstrap theory is capable of capturing that intuition. However, van Fraassen's conclusion may not be inescapable.

Plausible though it may appear, principle (P) has been denied by, among oth-

ers, Levi (1967), Kaplan (1981a, 1981b), Lehrer (1990), and Maher (1993), who have argued that, loosely, justification has the structure of a decision-making problem. In such a problem, one heeds not only the probabilities of the possible outcomes of a certain decision, but also their utilities. More exactly, in decision making one chooses the option that has the greatest expected utility of the available alternatives, where an option's expected utility is just the sum of the utilities of its various possible outcomes weighted by the probabilities of those outcomes.¹⁶ According to the aforementioned authors, there is nothing in the way decision theory is set up that would prevent applying it to matters epistemological; we can perfectly well assign cognitive or epistemic utilities to the 'acts' of accepting, rejecting, and suspending judgement on particular hypotheses or theories under particular circumstances, and then apply the decision-theoretic apparatus to these acts in the normal manner in order to determine what the agent is justified to do. So, on this approach to justification, a person may well be more justified in believing one theory than he is in believing a second even if the former is less likely to him than the latter, because, given his probabilities and utilities, the former may well have a greater expected cognitive utility than the latter.¹⁷

Now the notion of utility is anything but crystal-clear.¹⁸ The notion of cognitive utility appears even more problematic. We are told that the cognitive utility of accepting some hypothesis depends on the informativeness of that hypothesis (cf., e.g., Lehrer 1990, Maher 1993). But the notion of informativeness itself is still very much in need of clarification. Maher thinks this notion is to be cashed out in terms of verisimilitude. However, given that there is still widespread disagreement over the nature of verisimilitude (cf., e.g., Niiniluoto 1998), this suggestion

¹⁶See Jeffrey (1983) for a lucid presentation of the theory's basic machinery; also Resnik (1987).

¹⁷It would take me too far afield here to discuss van Fraassen's reasons for rejecting this view on justification. For criticisms of these reasons, see, e.g., Kukla (1998), Douven (1999, 2002, 2005c), Niiniluoto (1999), and Psillos (1999).

¹⁸As, e.g., Gillies (2000) and Howson (2000) and (2003) have argued (whether a defense of Bayesianism that does not appeal to utilities is possible, as Howson claims, is doubtful, however; cf. Douven (2003)). For one, it is entirely unclear whether an agent's risk-averseness should be reflected in his utility function; cf. Weirich (1986 and 2001), Rabin (2000), and Hacking (2001: 100f) for discussion. As a further indication of the unclarity, see the divergent interpretations of utility proposed in, for instance, Hansson (1988), Hampton (1994), and Dreier (1996). Some believe that utilities are just theoretical posits that do not stand in need of any interpretation (this view seems to underlie Ramsey's (1926) and Savage's (1954) work in decision theory and is still not uncommon, as Rabin (2000) reminds us). But aside from the difficulties generally related to instrumentalist interpretations of theoretical terms, on an instrumentalist reading decision theory, and hence also cognitive decision theory, can only be used as an explanatory device and not as a guide to decision making (see, e.g., Satz and Ferejohn 1994). So, in particular, cognitive decision theory could on that reading not inform us about when it is rational for us to believe a particular hypothesis or theory; at most it could be used post factum to explain why someone preferred to accept one rather than another hypothesis or theory.

seems rather unhelpful.¹⁹ One way in which my quantitative bootstrap account could be positively related to justification is by replacing, in cognitive decision theory, the ill-defined notion of cognitive utility by the clearly defined notion of degree of bootstrap support. In order to determine the justificational status of a hypothesis or theory we would thus have to weigh not probability and utility but probability and degree of bootstrap support against one another. This is only a rough proposal that can be filled out in quite diverse ways. I will not explore the possibilities here. My aim in this section merely was to point out that there may still be a positive role for degree of bootstrap confirmation even though a higher degree of bootstrap support does not generally indicate a higher probability.²⁰

5.6 B_m as Measures of the Coherence Between a Theory and the Evidence

The bootstrap measures of confirmation presented in this chapter bear some close resemblances to the measures C_m presented in chapter 2. Both explicate an intuitive concept (either that of coherence as mutual support or that of bootstrap confirmation) in terms of the mutual support relations between the different elements of the set or theory, and both measure support by means of an adaptation of some of the more popular Bayesian measures of support. However, as the last of the eight remarks in section 5.2 already suggested, it seems that the analogy can go further than that.

In the introductory chapter I briefly discussed Bonjour's challenge to foundationalists that they, too, require a theory of coherence. Recall that according to the foundationalist there are foundational beliefs that do not require any further justification; they are, so to speak, 'self-justifying'. Coherentists deny that there are beliefs that have such a special status. Instead they believe that all of our

¹⁹See Goossens (1976) for a more systematic critique of the concept of cognitive utility.

²⁰A way in which bootstrap confirmation could play a role in determining the justificational status of a theory that respects principle (P) is to assign a justificatory role to bootstrap support only after probabilistic considerations have been taken account of. This is, for instance, what the following principle does:

(P*) If (i) $p(T|E) > p(T'|E)$ or (ii) $p(T|E) = p(T'|E)$ and E bootstrap confirms T but not T' or (iii) $p(T|E) = p(T'|E)$ and E bootstrap confirms both T and T' but $B_{m^*}(T, E) > B_{m^*}(T', E)$ for one's preferred measure of bootstrap confirmation m^* , then we are more justified in believing T than we are in believing T' .

Again another response to the puzzle would be to claim that justification is to be evaluated at the level of single hypotheses and not at the level of theories. It is perfectly compatible with one theory as a whole being more probable than another theory as a whole that the probability of any of the axioms of the latter exceeds the probability of each of the axioms of the former. This would be much along the lines of Merricks's (1995) response to Klein and Warfield's (1994) claim that coherence is not generally truth conducive.

beliefs are justified only relative to other beliefs, the justificational status of any specific belief then being determined by how well it coheres with our other beliefs. Clearly, for such a theory to be successful we require a theory of coherence. As may be recalled from my discussion in chapter 4, I believe that the difference function presented there goes at least some way toward providing such a theory.

However, according to Bonjour, foundationalists, too, require a theory of coherence:

[T]he concept of coherence ... is also an indispensable ingredient in virtually all foundationalist theories; coherence must seemingly be invoked to account for the relation between the basic or foundational beliefs and other nonfoundational or 'superstructure' beliefs, in virtue of which the latter are justified in relation to the former. (Bonjour 1999: 124)

According to Bonjour (1999: 140n), 'strictly deductive or even enumerative inductive inference from the foundationalist beliefs does not suffice to justify most of the superstructure beliefs that the foundationalist typically wants to claim to be justified.' Without some recourse to the concept of coherence the foundationalist cannot justify the superstructure beliefs by reference to the foundational beliefs. As Douven (2005b) remarks, it seems likely that the notion of coherence that coherentists need to employ is different from the one required for foundationalism to succeed:

the foundationalist concept of coherence must pertain to a relation that holds between a body of nonbasic beliefs ... and a body of basic beliefs... That is, foundationalist coherence involves beliefs that do not, at least not initially, all have the same justificational status. For the coherentist, who denies that there is any kind of justification short of that provided by relations of coherence, coherence obtains or fails to obtain between beliefs that have, or at least may all have, the same justificational status. (Douven 2005b)

In his (2005b), Douven presents a theory of the coherence between the foundational (or evidential, in his terminology) and nonfoundational (or theoretical) beliefs, based on the measure of bootstrap confirmation \mathcal{B}_δ . Let me adopt the latter terminology.

In chapter 2 I remarked that we have two general intuitions with respect to coherence, viz., that it is a matter of how well the various propositions involved hang together and that it is a matter of degree. It seems clear that all three measures of bootstrap confirmation respect both of these intuitions. As noted above, the qualitative definition of bootstrap confirmation (definition 5.1) can be interpreted as indicating the coherence of the axioms of a theory with each other and the evidence. Moreover, now that I have defined a class of measures of bootstrap confirmation, the second intuitive desideratum is satisfied also.

This suggests that we can interpret the measures \mathcal{B}_m ($m \in \{\delta, r, c\}$) as measures of the coherence between a theory and the evidence. Such a strategy would

throw an interesting new light on some of the issues raised in this chapter. For example, if we follow the same reasoning as we did in chapter 2 then we can answer affirmatively the question of whether to add any weighing factors to the definition of \mathcal{B}_m : here, too, we could argue for a straight average as being the most plausible option. And if we would alter the definition in that way, the problem of independent propositions would no longer arise, for in that case it is easy to check that adding an independent proposition to a theory with a positive degree of bootstrap confirmation will always decrease the value of \mathcal{B}_m for all $m \in \{\delta, r, c\}$.

Seen from the perspective of the coherence between a theory and the evidence, the problem raised by van Fraassen acquires a new meaning also. For from that point of view, the question bears close resemblance to the discussion concerning the truth conduciveness of coherence. If my arguments in chapter 1 are correct, then the fact that the coherence between a theory and the evidence cannot be a truth conducive property *per se* does not affect the value of a measure of this type of coherence in any way.

5.7 Conclusion

This chapter has presented my proposal for an expression of bootstrap confirmation, which I have based on a few simple intuitive desiderata. One of the large projects still left open is to formulate other (possibly stricter) desiderata and see in what ways these would affect the different \mathcal{B}_m ($m \in \{\delta, r, c\}$). It would be interesting to see how robust the various theorems in this chapter are if alternative desiderata lead to different versions of \mathcal{B}_m or require us to adapt different measures of confirmation to measuring bootstrap confirmation. Moreover, other desiderata may help us to decide which of the measures of confirmation adapted to bootstrap confirmation is the most satisfactory. For whereas the different consequences discussed in section 5.4 do not seem to point one way or the other, it seems not unreasonable to suppose that other desiderata may be capable of singling out one of the three as the most satisfactory measure of bootstrap confirmation.

In contrast to the foregoing chapters, this chapter has not presented a measure of coherence but a measure of confirmation. However, in the last section I have described how an interpretation of the measure of bootstrap confirmation as a measure of the coherence between a theory and the evidence can be given shape. Moreover, it seems clear that such an interpretation will at least answer the problems noted in section 5.5. One of the problems that it will not solve, however, is the problem of different axiomatizations of a theory: in general different axiomatizations of a theory will have a different degree of bootstrap confirmation as measured by \mathcal{B}_m ($m \in \{\delta, r, c\}$). This problem will be taken up in the next

chapter.

Appendix 5 A: Proofs of the Theorems 5.1–5.7

Theorem 5.1 *For any hypothesis H , evidence E and complex of auxiliaries H' , it holds that $\delta(H; H', E)$, $\mathfrak{r}(H; H', E)$ and $\mathfrak{c}(H; H', E)$ are zero when H is independent of E and H' and positive (negative) if H is bootstrap confirmed (disconfirmed) by E with respect to H' .*

Proof: The part of the theorem about independence is clear enough. Next consider the case of (dis)confirmation for the adaptations of the difference measure and the log-ratio measure. First recognize that, by definition 5.3, if H is bootstrap confirmed by E relative to H' , then $p(H | H' \wedge E) > p(H | H')$ and $\delta(H; H', E)$ and $\mathfrak{r}(H; H', E)$ are both larger than zero. By definition 5.4, if H is bootstrap disconfirmed by E relative to H' , then $p(H | H' \wedge E) < p(H | H')$ and $\delta(H; H', E)$ and $\mathfrak{r}(H; H', E)$ are both smaller than zero.

Next consider the adaptation of Carnap's relevance measure \mathfrak{c} . First note that

$$p(H | H' \wedge E) = \frac{p(H \wedge H' \wedge E)}{p(H' \wedge E)}$$

Suppose the hypothesis is bootstrap confirmed by the evidence relative to the auxiliaries. Then:

$$\begin{aligned} & p(H | H' \wedge E) - p(H | H') > 0 \\ \Leftrightarrow & \frac{p(H \wedge H' \wedge E)}{p(H' \wedge E)} - p(H | H') > 0 \\ \Leftrightarrow & p(H \wedge H' \wedge E) - p(H | H')p(H' \wedge E) > 0 \\ \Leftrightarrow & \mathfrak{c} > 0 \end{aligned}$$

The case for disconfirmation follows directly from replacing ' $>$ ' by ' $<$ '. ■

Theorem 5.2 *Given a hypothesis H , evidence E and complex of auxiliaries H' it is not always the case that $\mathfrak{l}(H; H', E) > 0$ if H is bootstrap confirmed by E relative to H' .*

Proof: By the following probability model:

H	H'	E	probability	H	H'	E	probability
T	T	T	$7/512$	F	T	T	$15/128$
T	T	F	$3/64$	F	T	F	$1/2$
T	F	T	$1/4$	F	F	T	0
T	F	F	$1/32$	F	F	F	$21/512$

It can easily be checked that in this case $p(H | E \wedge H') \approx .105 > .0893 \approx p(H | H')$ and therefore, that H is bootstrap confirmed by E relative to H' . But also, $p(E \wedge H' | H) = .04 < .178 \approx p(E \wedge H' | \neg H)$ and thus $\mathfrak{f}(H; H', E) < 0$. ■

Theorem 5.3 *For all T and E , if E bootstrap confirms T , then $\mathcal{B}_m(T, E) > 0$ for all $m \in \{\delta, r, c\}$.*

Proof: Let $T = \{H_1, \dots, H_n\}$ and assume that E bootstrap confirms T . It follows from definition 5.5 and theorem 5.1 that, for all $H_i \in T$ and all $H_{ij}^T \in \wp(T \setminus \{H_i\})$, $m(H_i; H_{ij}^T, E) \geq 0$. Further, it also follows that for all $H_i \in T$ there is at least one $H_{ij}^T \in \wp(T \setminus \{H_i\})$ such that $m(H_i; H_{ij}^T, E) > 0$ for all $m \in \{\delta, r, c\}$. Since \mathcal{B}_m just adds up the bootstrap support each hypothesis receives relative to each set of possible auxiliaries as measured by $m \in \{\delta, r, c\}$, it follows that $\mathcal{B}_m(T, E) > 0$ for all $m \in \{\delta, r, c\}$. ■

Theorem 5.4 *For all T and E and for all $m \in \{\delta, r, c\}$, if $\mathcal{B}_m(T, E) < 0$, then E bootstrap disconfirms T .*

Proof: Let $T = \{H_1, \dots, H_n\}$. If $\mathcal{B}_m(T, E) < 0$, then, by definition 5.7, there must be at least one $H_i \in T$ such that, for at least one $H_{ij}^T \in \wp(T \setminus \{H_i\})$: $m(H_i; H_{ij}^T, E) < 0$. From definition 5.6 and theorem 5.1, it follows that in that case E bootstrap disconfirms T . ■

Theorem 5.5 *There is no $a \in \mathbb{R}$ and no $m \in \{\delta, r, c\}$ such that, for all T and E , if $\mathcal{B}_m(T, E) > a$, then E bootstrap confirms T .*

Proof: Toward a reductio, suppose that, for all T, E , if $\mathcal{B}_m(T, E) > c$ for some particular $c \in \mathbb{R}$ and one particular $m \in \{\delta, r, c\}$, then E bootstrap confirms T . Then let $T' = \{H_1, \dots, H_n\}$ and furthermore let it be the case that $\mathcal{B}_m(T', E') > c$ for some E' and some $m \in \{\delta, r, c\}$. Now let $T'' = \{H_1, \dots, H_n, H_{n+1}\}$ with H_{n+1} any hypothesis that is probabilistically independent of any subset of $\{H_1, \dots, H_n, E'\}$ (one can, without loss of generality, assume that such an H_{n+1} exist). Thus in particular the following facts hold:

- (i) $p(H_i | H_{ij}^{T'} \wedge H_{n+1}) = p(H_i | H_{ij}^{T'})$ for all $H_i \in T'$ and all $H_{ij}^{T'} \in \wp(T' \setminus \{H_i\})$;
- (ii) $p(H_i | H_{ij}^{T'} \wedge H_{n+1} \wedge E') = p(H_i | H_{ij}^{T'} \wedge E')$ for all $H_i \in T'$ and all $H_{ij}^{T'} \in \wp(T' \setminus \{H_i\})$;
- (iii) $p(E' \wedge H_{ij}^{T'} \wedge H_{n+1}) = p(E' \wedge H_{ij}^{T'})p(H_{n+1})$;
- (iv) $p(E' \wedge H_{ij}^{T'} \wedge H_{n+1} \wedge H_i) = p(E' \wedge H_{ij}^{T'} \wedge H_i)p(H_{n+1})$;

- (v) $p(H_{n+1} \mid H_{n+1k}^{T''} \wedge E') = p(H_{n+1} \mid H_{n+1k}^{T''}) = p(H_{n+1})$ for all $H_{n+1k}^{T''} \in \wp(T'' \setminus \{H_{n+1}\})$.

Dividing into three parts the sum that gives the value of $\mathcal{B}_m(T'', E')$ (see definition 5.7), yields

$$\begin{aligned} \mathcal{B}_m(T'', E') = & \sum_{i=1}^n \sum_{j=1}^{2^{n-1}} m(H_i; H_{i_j}^{T'}, E') + \\ & \sum_{i=1}^n \sum_{j=1}^{2^{n-1}} m(H_i; H_{i_j}^{T'} \wedge H_{n+1}, E') + \sum_{k=1}^{2^n} m(H_{n+1}; H_{n+1k}^{T''}, E'). \end{aligned} \quad (5.8)$$

Given (i) and (ii), we have, for all i such that $1 \leq i \leq n$ and all $H_{i_j}^{T'}$ such that $1 \leq j \leq 2^{n-1}$, that $p(H_i \mid H_{i_j}^{T'} \wedge E') - p(H_i \mid H_{i_j}^{T'}) = p(H_i \mid H_{i_j}^{T'} \wedge H_{n+1} \wedge E') - p(H_i \mid H_{i_j}^{T'} \wedge H_{n+1})$, and thus also that $\delta(H_i; H_{i_j}, E') = \delta(H_i; H_{i_j} \wedge H_{n+1}, E')$ and $r(H_i; H_{i_j}, E') = r(H_i; H_{i_j} \wedge H_{n+1}, E')$.

In similar fashion it follows from (i), (iii), and (iv) that $c(H_i; H_{i_j}, E') = c(H_i; H_{i_j} \wedge H_{n+1}, E')p(H_{n+1})$.

From this it follows that the first two summands in 5.8 are proportional to each other:

$$\sum_{i=1}^n \sum_{j=1}^{2^{n-1}} m(H_i; H_{i_j}^{T'}, E') = \alpha \left(\sum_{i=1}^n \sum_{j=1}^{2^{n-1}} m(H_i; H_{i_j}^{T'} \wedge H_{n+1}, E') \right), \quad (5.9)$$

where α equals either 1 (for the first two measures) or $p(H_{n+1})$ (for the bootstrap relevance measure), and therefore $\alpha \geq 0$.

Furthermore, from (v) together with theorem 5.1 it can be immediately seen to follow that

$$\sum_{k=1}^{2^n} m(H_{n+1}; H_{n+1k}^{T''}, E') = 0. \quad (5.10)$$

Since the first of the summands in 5.8 equals $\mathcal{B}_m(T', E')$ for all $m \in \{\delta, r, c\}$, it follows from 5.9 and 5.10 that $\mathcal{B}_m(T'', E') = (1 + \alpha)\mathcal{B}_m(T', E')$ for all $m \in \{\delta, r, c\}$. Since, furthermore, by supposition $\mathcal{B}_m(T', E') > c$, it must be the case that $\mathcal{B}_m(T'', E') > c$. Thus, E' bootstrap confirms T'' . However, it follows from equation (5.10) that clause 1 of definition 5.5 is not satisfied, so that E' does not bootstrap confirm T'' . Hence, the assumption that there is a numerical threshold for bootstrap confirmation leads to contradiction. Hence, there is no such numerical threshold. ■

Theorem 5.6 *There is no $a \in \mathbb{R}$ and no $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$ such that, for all T and E , if E bootstrap disconfirms T , then $\mathcal{B}_{\mathbf{m}}(T, E) < a$.*

Proof: First, if E bootstrap disconfirms T , then that does not exclude that $\mathcal{B}_{\mathbf{m}}(T, E) > 0$ for any $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$. To see this, consider a slightly changed version of the model given in example (5.1). Let in that model $p(E \wedge H_i \wedge H_j \wedge H_k) = .01$ (instead of .015), for all $i, j \in \{1, \dots, 4\}$ such that $i \neq j$. Then E bootstrap disconfirms T (for E disconfirms every $H_i \in T$ relative to the conjunction of every $H_j, H_k \in T$ such that $i \neq j \neq k$: $p(H_i | H_j \wedge H_k \wedge E) = .01/.05 = .2 < .25 = p(H_i | H_j \wedge H_k)$). Still, as an easy calculation shows, $\mathcal{B}_{\delta}(T, E) = 1.72$, $\mathcal{B}_{\mathbf{r}}(T, E) \approx 4.6372$ and $\mathcal{B}_{\mathbf{c}}(T, E) = 2.002$ and thus all $\mathcal{B}_{\mathbf{m}}(T, E) > 0$ are positive. Second, in the proof of theorem 5.5 it turned out that if one adds a hypothesis H to any theory T that is probabilistically independent of that theory together with the evidence E (in the precise sense specified in that proof), then $\mathcal{B}_{\delta}(T', E) > \mathcal{B}_{\delta}(T, E)$ for $T' = T \cup \{H\}$ and for $p(H) > 0$. So let then T be bootstrap disconfirmed by E and such that $\mathcal{B}_{\mathbf{m}}(T, E) > c > 0$ for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$. Adding a probabilistically independent hypothesis to T will result in a theory T' that is also bootstrap disconfirmed by E but for which $\mathcal{B}_{\mathbf{m}}(T', E) > \mathcal{B}_{\mathbf{m}}(T, E)$ for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$. Since this procedure can be repeated as often as one likes, there can be no $a \in \mathbb{R}$ such that $\mathcal{B}_{\mathbf{m}}(T, E) \geq a$ indicates that T is not bootstrap disconfirmed by E . ■

Theorem 5.7 *For all T, T' , and E , if $T \subset T'$ and E does not bootstrap disconfirm T' , then $\mathcal{B}_{\mathbf{m}}(T', E) \geq \mathcal{B}_{\mathbf{m}}(T, E)$ for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$; if in addition E bootstrap confirms T' , then $\mathcal{B}_{\mathbf{m}}(T', E) > \mathcal{B}_{\mathbf{m}}(T, E)$ for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$.*

Proof: Let $T' = \{H_1, \dots, H_n\}$. Without loss of generality, assume that $T = \{H_1, \dots, H_m\}$ ($m < n$). From definition 5.7 it follows that for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$

$$\begin{aligned} \mathcal{B}_{\mathbf{m}}(T', E) &= \sum_{i=1}^n \sum_{j=1}^{2^{n-1}} \mathbf{m}(H_i; H_{i_j}^{T'}, E) \\ &= \sum_{i=1}^m \sum_{j=1}^{2^{n-1}} \mathbf{m}(H_i; H_{i_j}^{T'}, E) + \sum_{k=m+1}^n \sum_{l=1}^{2^{n-1}} \mathbf{m}(H_k; H_{k_l}^{T'}, E). \end{aligned}$$

Now note that, if E does not bootstrap disconfirm T' , it must hold for all $H_i \in T$ and all $H_{i_j}^{T'} \in \wp(T' \setminus \{H_i\})$ that $\mathbf{m}(H_i; H_{i_j}^{T'}, E) \geq 0$. Thus, for all $\mathbf{m} \in \{\delta, \mathbf{r}, \mathbf{c}\}$

$$\mathcal{B}_{\mathbf{m}}(T, E) = \sum_{i=1}^m \sum_{j=1}^{2^{m-1}} \mathbf{m}(H_i; H_{i_j}^T, E) \leq \sum_{i=1}^m \sum_{j=1}^{2^{n-1}} \mathbf{m}(H_i; H_{i_j}^{T'}, E).$$

Hence, if $\sum_{k=m+1}^n \sum_{l=1}^{2^{n-1}} \mathfrak{m}(H_k; H_{k_l}^{T'}, E) \geq 0$, the following must hold for all $\mathfrak{m} \in \{\mathfrak{d}, \mathfrak{r}, \mathfrak{c}\}$:

$$\mathcal{B}_{\mathfrak{m}}(T', E) = \sum_{i=1}^n \sum_{j=1}^{2^{n-1}} \mathfrak{m}(H_i; H_{i_j}^{T'}, E) \geq \sum_{i=1}^m \sum_{j=1}^{2^{m-1}} \mathfrak{m}(H_i; H_{i_j}^T, E) = \mathcal{B}_{\mathfrak{m}}(T, E). \quad (5.11)$$

But it is easy to see that the condition is satisfied. For since T' is not bootstrap disconfirmed by E , it must be the case for all $H_i \in T' - T$ that for all $H_{i_j}^{T'} \in \wp(T' \setminus \{H_i\})$, we have $\mathfrak{m}(H_i; H_{i_j}^{T'}, E) \geq 0$ (in virtue of theorem 5.1). If E bootstrap confirms T' , then it must also be the case that for at least one $H_{i_j}^{T'} \in \wp(T' \setminus \{H_i\})$, we have $\mathfrak{m}(H_i; H_{i_j}^{T'}, E) > 0$ (by definition 5.5 and theorem 5.1) so that

$$\sum_{k=m+1}^n \sum_{l=1}^{2^{n-1}} \mathfrak{m}(H_k; H_{k_l}^{T'}, E) > 0,$$

whence it follows that in equation (5.11) the ' \geq ' can be replaced by a '>'. ■

Appendix 5 B: Proofs of the Theorems 5.8 and 5.9

Theorem 5.8 *For all $n \in \mathbb{N}$, T , and E , if $T = \{H_1, \dots, H_n\}$ and $\mathcal{B}_m(T, E) \geq n 2^{n-1} - 1$ for $m \in \{\delta, c\}$, then E bootstrap confirms T .*

Proof: Suppose $\mathcal{B}_m(T, E) \geq n 2^{n-1} - 1$ for all $m \in \{\delta, c\}$ and for some E and $T = \{H_1, \dots, H_n\}$. Then it follows from definition 5.7 that for no $H_i \in T$ can it be the case that there is a $H_{ij}^T \in \wp(T \setminus \{H_i\})$ such that $m(H_i; H_{ij}^T, E) \leq 0$ with $m \in \{\delta, c\}$. Hence for all $H_i \in T$ and all $H_{ij}^T \in \wp(T \setminus \{H_i\})$ it must hold that $p(H_i | H_{ij}^T \wedge E) > p(H_i | H_{ij}^T)$. And thus, by definition 5.5, E bootstrap confirms T . (To see that, if for even a single H_i the value of $m(H_i; H_{ij}^T, E)$ with $m \in \{\delta, c\}$ is lower than or equal to 0 for some set of auxiliaries, then the value of $\mathcal{B}_m(T, E)$ must be strictly smaller than, and hence cannot be equal to, $n 2^{n-1} - 1$, one only has to note that the range of both δ and c is the open interval $(-1, 1)$, and that -1 respectively 1 are not within the range because (1) for $p(H | H' \wedge E) - p(H | H')$ to obtain those values, it would have to hold that $p(H | H' \wedge E) = 0$ and at the same time that $p(H | H') = 1$, respectively, that $p(H | H' \wedge E) = 1$ and at the same time that $p(H | H') = 0$, neither of which combinations is possible and (2) for $p(H \wedge H' \wedge E) - p(H | H')p(H' \wedge E)$ to obtain those values, it would have to hold that $p(H \wedge H' \wedge E) = 0$ and at the same time that $p(H | H') = p(H' \wedge E) = 1$, respectively, that $p(H \wedge H' \wedge E) = 1$ and at the same time that $p(H | H') = 0$ or $p(H' \wedge E) = 0$, again neither of which combinations is possible.) ■

Theorem 5.9 *For all $n \in \mathbb{N}$, T , and E , if $T = \{H_1, \dots, H_n\}$, $\mathcal{B}_r(T, E)$ can be arbitrarily large while E still bootstrap disconfirms T .*

Proof: Consider a theory $T = \{H_1, \dots, H_n\}$ which is bootstrap disconfirmed by the evidence E . Then for at least one H_i and one $H_{ij}^T \in \wp(T \setminus \{H_i\})$, $p(H_i | H_{ij}^T \wedge E) < p(H_i | H_{ij}^T)$. Now suppose for another hypothesis H_k ($k \neq i$) and some subset $H_{k_l}^T \in \wp(T \setminus \{H_k\})$ that $p(H_k | H_{k_l}^T)$ approaches, but is not equal to, zero, while $p(H_k | H_{k_l}^T \wedge E)$ approaches, but is not equal to, unity. It can easily be seen that probability models to that effect exist. But in that case, the closer $p(H_k | H_{k_l}^T)$ is to zero, the larger $r(H_k; H_{k_l}^T, E) = \log[p(H_k | H_{k_l}^T \wedge E)/p(H_k | H_{k_l}^T)]$ becomes. In the limit, $r(H_k; H_{k_l}^T, E)$ will go to infinity and therefore $\mathcal{B}_r(T, E)$ will likewise go to infinity. Thus, $\mathcal{B}_r(T, E)$ can be arbitrarily large even if T is bootstrap disconfirmed. Thus, there is no numerical threshold for $\mathcal{B}_r(T, E)$ above which T is always bootstrap confirmed. ■

5 C: Proofs of Theorems 5.10–5.14

Theorem 5.10 *For all T and E , if E bootstrap confirms T , then E also probabilistically confirms T .*

Proof: Let $T = \{H_1, \dots, H_n\}$ and suppose E bootstrap confirms T . Then, by clause 1 of definition 5.5, there must for every $H_i \in T$ be at least one $T' \subset T \setminus \{H_i\}$ such that $p(H_i | E \wedge T') > p(H_i | T')$. Thus in particular there must for H_1 be a subset T^* of $T \setminus \{H_1\}$ such that

$$p(H_1 | E \wedge T^*) > p(H_1 | T^*). \quad (5.12)$$

Now let $\pi_1, \dots, \pi_n!$ denote the permutations on $1, \dots, n$. Clearly,

$$\begin{aligned} p(E \wedge T) &= p(E)p(H_{\pi_1(1)} | E) \cdots p(H_{\pi_i(n)} | E \wedge H_{\pi_i(1)} \wedge \cdots \wedge H_{\pi_i(n-1)}) \\ &= p(E)p(H_{\pi_j(1)} | E) \cdots p(H_{\pi_j(n)} | E \wedge H_{\pi_j(1)} \wedge \cdots \wedge H_{\pi_j(n-1)}), \end{aligned}$$

for all $i, j \in \{1, \dots, n!\}$. Observe that for some $k \in \{1, \dots, n!\}$, $p(H_1 | E \wedge T^*)$ must occur as a factor in $p(E)p(H_{\pi_k(1)} | E) \cdots p(H_{\pi_k(n)} | E \wedge H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(n-1)})$. Next suppose, toward a reductio, that the consequent of the theorem does not hold, i.e., E does not probabilistically confirm T . Then $p(T | E) \leq p(T)$, or $[p(T \wedge E) / p(E)] \leq p(T)$, or again written differently,

$$\frac{p(E \wedge H_1 \wedge \cdots \wedge H_n)}{p(E)} \leq p(H_1 \wedge \cdots \wedge H_n).$$

By the general multiplication rule and after canceling $p(E)$ in the left-hand expression, this is equivalent to

$$\begin{aligned} p(H_1 | E) \cdots p(H_n | E \wedge H_1 \wedge \cdots \wedge H_{n-1}) &\leq \\ p(H_1)p(H_2 | H_1) \cdots p(H_n | H_1 \wedge \cdots \wedge H_{n-1}). \end{aligned} \quad (5.13)$$

Given (5.13), the following must also hold:

$$\begin{aligned} p(H_{\pi_k(1)} | E) \cdots p(H_{\pi_k(n)} | E \wedge H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(n-1)}) &\leq \\ p(H_{\pi_k(1)})p(H_{\pi_k(2)} | H_{\pi_k(1)}) \cdots p(H_{\pi_k(n)} | H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(n-1)}). \end{aligned} \quad (5.14)$$

From inequality (5.12) we know that for one i with $1 \leq i \leq n$, it must hold that

$$p(H_{\pi_k(i)} | E \wedge H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(i-1)}) > p(H_{\pi_k(i)} | H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(i-1)}).$$

Combining this with (5.14), we get that for at least one j with $1 \leq j \leq n$,

$$p(H_{\pi_k(j)} | E \wedge H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(j-1)}) < p(H_{\pi_k(j)} | H_{\pi_k(1)} \wedge \cdots \wedge H_{\pi_k(j-1)}),$$

for else the left-hand side of (5.14) will be larger than the right-hand side. It follows that there is a $H_i \in T$ and a $T' \subset T$ such that $p(H_i | E \wedge T') < p(H_i | T')$. But this violates clause 2 of definition 5.5 and hence our assumption that E bootstrap confirms T . Thus the assumption that E does not probabilistically confirm T is false. ■

Theorem 5.11 *For all T and E , if E probabilistically disconfirms T , then E also bootstrap disconfirms T .*

Proof: Assume the antecedent, i.e., $p(T | E) < p(T)$, where $T = \{H_1, \dots, H_n\}$. Then

$$\frac{p(T \wedge E)}{p(E)} < p(T),$$

or, with T written out,

$$\frac{p(E \wedge H_1 \wedge \dots \wedge H_n)}{p(E)} < p(H_1 \wedge \dots \wedge H_n).$$

Multiplying both sides by $p(E)$ yields

$$p(E \wedge H_1 \wedge \dots \wedge H_n) < p(E)p(H_1 \wedge \dots \wedge H_n).$$

Using the general multiplication rule for both sides, we obtain

$$\begin{aligned} p(E)p(H_1 | E) \dots p(H_n | E \wedge H_1 \wedge \dots \wedge H_{n-1}) < \\ p(E)p(H_1) \dots p(H_n | H_1 \wedge \dots \wedge H_{n-1}), \end{aligned}$$

which can only be the case if

$$\begin{aligned} [p(H_1 | E) < p(H_1)] \vee \dots \vee \\ [p(H_n | E \wedge H_1 \wedge \dots \wedge H_{n-1}) < p(H_n | H_1 \wedge \dots \wedge H_{n-1})]. \end{aligned}$$

Thus there is at least one $H_i \in T$ and at least one $T' \subset T \setminus \{H_i\}$ such that $p(H_i | E \wedge T') < p(H_i | T')$ and hence, by definition 5.6, E bootstrap disconfirms T . ■

Theorem 5.12 *It is not the case that, for all T and E , if E probabilistically confirms T , then also E bootstrap confirms T .*

Proof: Consider the probability model obtained in example 5.2 the last clause saying E does not bootstrap disconfirm any of the $H_i \in T$ relative to any of the

remaining auxiliaries is omitted and one lets $a = .003$ and $x = .004$. In this case, E does not bootstrap confirm T , since for all $i, j \in \{1, 2, 3\}$:

$$p(H_i | E \wedge H_j) = \frac{p(E \wedge H_i \wedge H_j)}{p(E \wedge H_j)} = \frac{.004}{.05} = .08,$$

which is smaller than $.15 (= p(H_i | H_j))$. A fortiori, this model violates clause 2 of definition 5.5. However, E does probabilistically confirm T , for

$$p(T | E) = \frac{p(E \wedge H_1 \wedge H_2 \wedge H_3)}{p(E)} = \frac{.003}{.5} = .006 > .005 = p(T).$$

■

Theorem 5.13 *It is not the case that, for all T and E , if E bootstrap disconfirms T , then also E probabilistically disconfirms T .*

Proof: From the probability model constructed in the proof of theorem 5.12. ■

Theorem 5.14 *For all $m \in \{\delta, r, c\}$ there is no function f such that, for all T and E , $\mathcal{B}_m(T, E) = f \circ m(T, E)$, with m any non-bootstrap measure of confirmation.*

Proof: The proof of this theorem proceeds by specifying a probability model involving a theory T^* and evidence E^* for which the following hold (T and E are as in example 5.1): (i) $p(T^*) = p(T)$, (ii) $p(E^*) = p(E)$, (iii) $p(T^* | E^*) = p(T | E)$, (and thus also) (iv) $p(E^* | T^*) = p(E | T)$, (v) $p(T^* | \neg E^*) = p(T | \neg E)$, and (vi) $p(E^* | \neg T^*) = p(E | \neg T)$, but (vii) $\mathcal{B}_m(T^*, E^*) \neq \mathcal{B}_m(T, E)$ for all $m \in \{\delta, r, c\}$. It can readily be seen that, given (i)–(vi), and given how I defined the notion of a non-bootstrap measure of confirmation, there can be no function f such that $f \circ m(T, E) = \mathcal{B}_m(T, E)$ for all $m \in \{\delta, r, c\}$, where m is such a non-bootstrap measure of confirmation.

Like T , the theory T^* consists of four axioms, H_1^*, \dots, H_4^* . Like the hypotheses in T , the H_i^* all have a prior probability of .25, and are all mutually probabilistically independent; we thus see immediately that (i) holds. Evidence E^* has a prior probability of .5 (like E in example 5.1; so (ii) holds). Further we have the following:

- $p(E^* \wedge H_i^*) = .125$, for all $i \in \{1, \dots, 4\}$;
- $p(E^* \wedge H_i^* \wedge H_j^*) = .045$, for all $i, j \in \{1, \dots, 4\}$ such that $i \neq j$;
- $p(E^* \wedge H_i^* \wedge H_j^* \wedge H_k^*) = .0125$, for all $i, j, k \in \{1, \dots, 4\}$ such that $i \neq j \neq k$;
- $p(E^* \wedge H_1^* \wedge \dots \wedge H_4^*) = .0038$.

From the fact that $p(E^* \wedge H_1^* \wedge \cdots \wedge H_4^*) = p(E \wedge H_1 \wedge \cdots \wedge H_4)$ it follows that $p(T^* | E^*) = p(T | E)$ (so (iii) holds; and given (i)–(iii), (iv) and, by the law of total probability, (v) and (vi) must hold as well). Now calculate $\mathcal{B}_8(T^*, E^*)$ from the following values, which hold for all $i \in \{1, \dots, 4\}$:

- $\delta(H_i^*; E^*) = .125/.5 - .25 = 0$;
- $\delta(H_i^*; H_j^*, E^*) = .045/.125 - .25 = .11$, for all $j \in \{1, \dots, 4\}$ such that $i \neq j$;
- $\delta(H_i^*; H_j^* \wedge H_k^*, E^*) = .0125/.045 - .25 \approx .0278$, for all $j, k \in \{1, \dots, 4\}$ such that $i \neq j \neq k$;
- $\delta(H_i^*; H_j^* \wedge H_k^* \wedge H_l^*, E^*) = .0038/.0125 - .25 = .054$, for all $j, k, l \in \{1, \dots, 4\}$ such that $i \neq j \neq k \neq l$.

So, the bootstrap support each of the H_i^* gets from E^* as measured by \mathcal{B}_8 totals (approximately): $(1)(0) + (3)(.11) + (3)(.0278) + (1)(.054) = .4674$. And thus $\mathcal{B}_8(T^*, E^*) \approx (4)(.4674) = 1.8696$. This is unequal to the bootstrap support T was seen to get from E , namely 2.412, despite the fact that, as we saw, (i)–(vi) hold, and thus on any non-bootstrap measure of confirmation m , we have $m(T^*, E^*) = m(T, E)$.

In similar fashion one calculates that $\mathcal{B}_r(T^*, E^*) \approx 6.4223$, while $\mathcal{B}_r(T, E) \approx 7.8809$ and that $\mathcal{B}_c(T^*, E^*) \approx 0.1827$, while $\mathcal{B}_c(T, E) \approx 0.2552$. ■

Chapter 6

Anticipated Objections

In this chapter I will discuss two recent general arguments against the project of measuring coherence probabilistically: one by Siebel (2005) and the other one by Luca Moretti and Ken Akiba (2005). Although Moretti and Akiba do not make this claim themselves, it will become clear that if any of these two arguments is correct, there can be no satisfactory probabilistic measures of coherence. Nonetheless, my conclusion will be that at most these criticisms succeed at limiting the intended scope of such measures. For it will appear that if we exclude from our intended domain sets or theories in which the propositions or axioms are logically related, the problems no longer arise.

The objections discussed in this chapter are objections that can be raised against all of the measures of coherence proposed in this thesis (and the second objection even applies to the measure of bootstrap confirmation as well). This does not imply that these are the only possible objections against these measures. Rather, each measure runs into some potential difficulties of its own. Most notably, there is not one measure that can handle perfectly all the ‘test-cases’ proposed in the literature. For example, measures of mutual support will not satisfy the examples based on intuitions of relative overlap, while relative overlap measures perform poorly with respect to examples in which our mutual support intuitions are pivotal. Such objections have all been discussed in the relevant chapters, and I will not rehearse them here. But I believe that the discussions in those chapters have made it clear that none of those objections can survive a critical discussion.

Below I will not discuss Olsson’s (2005) paper called ‘The Impossibility of Coherence.’ The main reason for this is that although this paper claims to show that coherence is ‘not definable,’ it clearly only concerns the sense of coherence as a confidence boosting property: Olsson explicitly takes the model of witnesses

reporting on propositions for granted and even uses a reliability parameter to model the reliability of the witnesses. Therefore, even if his argument is correct, it would still only apply to the explication of coherence as a confidence boosting property.¹

6.1 Siebel's Impossibility Argument

In his 'Probabilistic Measures of Coherence' (2005), Siebel asks us to consider an example in which a physicist is not sure anymore which value the voltmeter read in a certain experimental set-up. Consider the sets $S = \{A_1, A_2\}$ and $S' = \{A_1, A_3\}$, with

- A_1 : The voltage was 1 V;
- A_2 : The voltage was 1 V or 2 V;
- A_3 : The voltage was 1 V or 50 V.

According to Siebel the second set is much less coherent than the first:

It would ... appear that A_2 fits together with A_1 more than A_3 does because A_2 's alternative is much closer to A_1 than A_3 's alternative. If A_1 is false because the voltage is 2 V, we are at least in close neighbourhood to A_1 . But if A_1 is false because the voltage is 50 V, then the truth is miles away from A_1 . [A] purely probabilistic approach overlooks that coherence is also sensitive to such distances of numerical values, thereby ignoring an aspect which is highly important for scientists. (Siebel 2005)

It is easy to see that since A_1 implies both A_2 and A_3 , then if A_2 and A_3 have the same marginal probability, the probability models for S and S' are equal. Therefore, each purely probabilistic measure of coherence must judge both sets as equally coherent under these circumstances. According to the quote above, this is highly counterintuitive, due to the difference in numerical distance between 1 V and 2 V and 1 V and 50 V, respectively.

I disagree. Firstly, I am not sure what the concept of 'distance of numerical values' refers to. The difference between two values is always relative to the scale that has been chosen to represent the values with. For example, choosing logarithmic scales leads to the values of .01 and 1 being further apart than 1

¹This does not mean that I believe that Olsson's argument, considered as a criticism of explicating coherence as a confidence boosting property, is correct. Similarly to Bovens and Hartmann's (2003a: 19–22) impossibility result, his arguments only show that coherence cannot be a complete ordering. However, as I argued in chapter 4, this need be no impediment to explicating the sense of coherence as a confidence boosting property with the help of the notion of a quasi-ordering. It is not clear to me how Olsson would respond to such a challenge. Since his paper was written before the publication of Bovens and Hartmann's (2003a), the idea of the relation 'is more coherent than' being a quasi-ordering is not discussed in Olsson's paper.

and 50. More generally, it is always possible to introduce scales such that the distance between 1 and 50 is arbitrarily small. Therefore, the argument that 1 and 50 are very far apart presupposes an *a priori* argument for a specific scale – an argument which no physicist would ever endorse.

But secondly, and more importantly, I believe that Siebel chooses a probability model that is inconsistent with the driving intuition behind his argument. The reason a physicist (or anyone else, for that matter) would intuitively feel that the first set is more coherent than the second is that the outcome of 1 V is much more probable given a reading of 2 V than of 50 V. That is, intuitively, the propositions in the first set overlap much more than those in the second set do.

Everybody who has ever conducted a physical experiment in less than ideal settings knows that getting close is often as good as it gets. If standard physical theory predicts a value of 1.0 V, then one is often quite happy with a 1.2 V or even a 2.1 V result, depending on the cumulative measurement errors in the different pieces of equipment that have been used to conduct the experiment. However, this intuition can play no role in Siebel's model, since in this case the outcomes of 2 V and 50 V are equally likely given that the real value is 1 V. What actual experiment can we think of that would satisfy such a constraint? Well, it appears that it must be an experiment in which the outcome does not have a measurement error: only then is it possible that the 2 V and the 50 V outcomes are equally likely (otherwise it will always be the case that 2 V is more likely than 50 V, conditional on the value being 1 V).

But if there is no measurement error then either the probabilities of a reading of 2 V and a reading of 50 V have zero probability (in which case both sets simply consist of identical propositions so that there is no reason to argue that the coherence should be different) or the readings can only occur if something is wrong in the set-up of the experiment: maybe some pieces are broken or wrongly connected. Moreover, it is apparently assumed that on the condition of a faulty set-up the outcomes of 2 V and 50 V are equally likely. But in that case there is no reason to argue that the first set is more coherent than the second: both 2 V and 50 V point to a failure in the experimental set-up and from that point of view their closeness to 1 V is irrelevant. Siebel's intuitions are only plausible on the hypothesis that given a true value of 1 V, an outcome of 2 V is much more likely than an outcome of 50 V. But this is excluded by his probability model.

Therefore, it seems safe to conclude that on any of the possible interpretations of Siebel's example, his claim that we intuitively feel there is a difference in coherence while there is in fact no difference, is no longer true. If the different values can only occur in the case of malfunctioning equipment or if they are not possible at all, then I believe we intuitively would not feel there should be a difference in coherence between the two cases. If, on the other hand, the different possible outcomes could be the result of measurement errors, then clearly the

probability models for the two sets cannot be equal, since one outcome (2 V) is much more likely than the other (50 V).

6.2 Moretti and Akiba's Equi-coherence Principle

In their 'Probabilistic Measures of Coherence and the Problem of Belief Individuation' (2005), Moretti and Akiba show that all the measures that had been published at the time of their writing fail to satisfy, among other things, the *Equi-coherence Principle* (henceforth EC):

EC If S and S' are logically equivalent sets of beliefs, S and S' have the same degree of coherence,

where two sets S and S' are logically equivalent iff, for any $R \in S$, we have $S' \vdash R$ and for any $R' \in S'$, we have $S \vdash R'$ (cf. Moretti and Akiba 2005). If EC is not satisfied, it follows that 'there can be different sets of beliefs, which represent the same information set, that have different degrees of coherence' (*ibid.*). Moretti and Akiba call this the problem of *belief individuation*. It is not hard to see that all of the measures of coherence proposed in chapters 2–4 violate EC (I will discuss one example below). Moreover, a similar problem arises for the measure of bootstrap confirmation introduced in chapter 5. As the following example shows, two equivalent theories may be bootstrap confirmed to a different degree:²

Example 6.1 Consider again example (5.1). But now replace the original theory $T = \{H_1, H_2, H_3, H_4\}$ with $T^* = \{H_1, H_2, H_3 \wedge H_4\}$. Clearly $T \equiv T^*$ (T^* is just a different axiomatization of T) and if we accept the argument above it should be the case that $\mathcal{B}_m(T^*, E) = \mathcal{B}_m(T, E)$ for all $m \in \{\delta, r, c\}$. But it is easy to appreciate that this is not the case for any of the \mathcal{B}_m ($m \in \{\delta, r, c\}$). For instance, the degree of bootstrap support that T^* receives from the evidence as measured by the bootstrap difference measure equals $\mathcal{B}_\delta(T^*, E) \approx .572$, while $\mathcal{B}_\delta \approx 2.412$ (and similarly for the other measures). ♦

At first sight, EC may be quite plausible. For consider the following example. Suppose a murder suspect has so far told a rather incoherent story about his whereabouts the night the murder was committed. How could he make his story more coherent? Well, he might be able to boost the coherence of his defense

²Note, however, that in note 6 in chapter 5, I made the measure of bootstrap confirmation apply only to theories that are naturally axiomatized. It is not clear that T and T^* in the example below are in fact naturally axiomatized. For instance, if Gemes's notion of a natural axiomatization is considered and one accepts his (1993: 483n3) proposal for an additional condition, then T^* is not naturally axiomatized.

by adding to it some statement or statements that explain(s) away the apparent tension between the elements of his story so far. But instead he continues, 'In other words, ...,' and he tells the same story again, though, as announced, in other words. Next he says, 'To put it a third way, ...,' and he puts it a third way. Could it be that, just by rephrasing what he has already said, i.e., by adding to his defense some statements that (let us assume) are logically equivalent to ones already made, that defense comes to be coherent after all? Surely the answer must be negative. To increase the degree of coherence of his defense it seems as useless to rephrase what he has already said as it is to repeat it verbatim.

Much the same would seem to hold if the defendant were to go on by drawing inferences from what he has already testified. For instance, suppose that, having said that he wasn't in London on the day the murder was committed, but that he spent that whole day in Brussels, he goes on to say that he wasn't in Amsterdam that day, and that he wasn't in Paris, and that he wasn't in Berlin. Again, it is hard to see how this could do anything to increase the coherence of his defense; the defendant is just asserting things which, though he did not expressly assert them, he did already imply by what he asserted.

The foregoing at least suggests that there ought to be no difference between the degree of coherence of a set S and that of another set S' if $S \subset S'$ and for every proposition $R \in S' \setminus S$ there is an $R' \in S$ such that $R' \vdash R$. However, all the measures proposed in this thesis fail to satisfy this intuitive desideratum. To see this, consider a set $S = \{R_1, R_2, R_3, R_4\}$ such that $R_1 \equiv R_3$ and $R_1 \vdash R_4$, and a probability distribution $p(\cdot)$ on this set such that

$$\begin{aligned} p(R_1 \wedge R_2 \wedge R_3 \wedge R_4) &= .01 & p(R_1 \wedge \neg R_2 \wedge R_3 \wedge R_4) &= .2 \\ p(\neg R_1 \wedge R_2 \wedge \neg R_3 \wedge R_4) &= .05 & p(\neg R_1 \wedge R_2 \wedge \neg R_3 \wedge \neg R_4) &= .15 \\ p(\neg R_1 \wedge \neg R_2 \wedge \neg R_3 \wedge R_4) &= .1 & p(\neg R_1 \wedge \neg R_2 \wedge \neg R_3 \wedge \neg R_4) &= .49 \end{aligned}$$

Now consider sets $S^\alpha = \{R_1, R_2\}$, $S^\beta = \{R_1, R_2, R_3\}$, and $S^\gamma = \{R_1, R_2, R_4\}$. As the following table shows, these sets have different degrees of coherence on all of the measures of coherence (all values are rounded to three decimal places, 'ind' stands for indeterminate):

	C_d	C_r	C_l	\mathcal{R}	S	\mathcal{O}	\mathcal{F}
S^α	-.162	.227	.188	.024	.227	.024	-.683
S^β	-.012	2.494	ind	.357	1.08	.024	-.087
S^γ	.071	7.814	ind	.185	.63	.02	-.036

This means that both adding a proposition to S^α that is logically equivalent to R_1 and adding one that is logically implied by R_1 lead to sets with degrees of coherence different from that of the original set.

It is worth briefly noting that Olsson's measure comes closest to satisfying the intuitive desideratum that a measure of coherence be insensitive to the operation

of adding to a set (or subtracting from it) propositions logically implied by or even logically equivalent to ones already in the set. It is easy to see that on this measure it holds quite generally that adding logically equivalent propositions does not affect the degree of coherence. After all, because both $R_1 \wedge \dots \wedge R_n \equiv R_1 \wedge \dots \wedge R_n \wedge R_{n+1}$ and $R_1 \vee \dots \vee R_n \equiv R_1 \vee \dots \vee R_n \vee R_{n+1}$ whenever $R_{n+1} \equiv R_i$ for some $i \in \{1, \dots, n\}$, adding a logically equivalent proposition neither affects the numerator nor the denominator in the formula of Olsson's measure. Note, though, that it does *not* generally hold on this measure that adding logical consequences cannot affect the degree of coherence. Of course, if for some $i \in \{1, \dots, n\}$, $R_i \vdash R_{n+1}$, then again $R_1 \wedge \dots \wedge R_n \equiv R_1 \wedge \dots \wedge R_n \wedge R_{n+1}$. But the same condition does not guarantee that $R_1 \vee \dots \vee R_n \equiv R_1 \vee \dots \vee R_n \vee R_{n+1}$. Because $R_1 \vee \dots \vee R_{n+1}$ can be logically weaker, but not stronger, than $R_1 \vee \dots \vee R_n$, adding a logical consequence can affect coherence only negatively. That is to say, adding a logical consequence cannot raise the degree of coherence on Olsson's measure, but it can *lower* that degree – which is still to violate the desideratum.

It is further worth noting that Bovens and Hartmann's difference function does not satisfy the desideratum either, since it turns out that according to their theory both $\langle S^\alpha, S^\beta \rangle$ and $\langle S^\alpha, S^\gamma \rangle$ constitute indeterminate cases. And according to the revised version of Bovens and Hartmann's difference function proposed in chapter 4, f_r^i , it is the case that both S^β and S^γ are more *i*-coherent than S^α .³

From the above example it also follows that all known measures of coherence violate EC as well, except Olsson's measure. The remark above that Olsson's measure satisfies the desideratum that adding logically equivalent propositions does not affect the degree of coherence, may lead one to expect that it also satisfies EC. However, this is not the case. Instead, a quite simply argument seems to show that if EC is correct, then it will make *all* probabilistic measures of coherence redundant. For suppose that EC is correct, and compare a set $S = \{R_1, \dots, R_n\}$ with a set $S_c = \{R_c\}$ with R_c equivalent to the conjunction of all the propositions in set S . Then by EC, S and S_c must have the same degree of coherence. However, a set consisting of one proposition has only one probabilistic property, to wit, the probability of that proposition.⁴ Thus, if a probabilistic measure of coher-

³Nonetheless, this does not imply that on these accounts the coherence of the defendant's story in the example changes when he repeats logically equivalent versions of his story (see below).

⁴One may feel that this is too quick, for it appears that if the proposition R_c is composed of a number of different conjuncts, then the probabilities of each of the conjuncts may equally be considered to be part of the probabilistic features of the set $S_c = \{R_c\}$. However, this is no impediment to the argument, which only attempts to show that no satisfactory probabilistic measures of coherence can ever satisfy EC. And it must be clear that not all propositions R_c that are equivalent to the conjunction of a number of propositions R_1, \dots, R_n are indeed composed of a number of conjuncts. For example, in the Tweety example discussed in section 3.5, the proposition that Tweety is a penguin is equivalent to the conjunction of the propositions that Tweety is a bird and that Tweety cannot fly (relative to the background knowledge), but clearly it is not composed of a number of different conjuncts.

ence is to satisfy EC, then it must be a function only of the probability that the conjunction of the elements of a set S is true: $p(R_1 \wedge \dots \wedge R_n) = p(R_c)$. But in that case coherence would be nothing more than the standard definition of the probability of a set. Thus, a full acceptance of EC would make any measure of coherence redundant. But need we accept EC? And if we need not, need we accept the intuitive desideratum, which would still discredit all the known measures of coherence?

It seems that the answer depends partly on which sense of coherence we are trying to explicate. Firstly, EC seems to make little sense for Bovens and Hartmann's model: if coherence is to be a confidence boosting property and if the only parameters that determine our confidence are witness reliability, prior probability and coherence, then it follows mathematically that sets S and S_c must in some cases have a different degree of coherence. Indeed, an easy calculation shows that for not fully unreliable witnesses:

$$\frac{p(R_1 \wedge \dots \wedge R_n | \text{REPR}_1 \wedge \text{REPR}_2 \wedge \dots \wedge \text{REPR}_n)}{p(R_1 \wedge \dots \wedge R_n | \text{REP}[R_1 \wedge R_2 \wedge \dots \wedge R_n])} > \quad (6.1)$$

Since by definition the prior probabilities of both sets are equal and since there are certainly cases in which the reliability is fixed also (for example, if the witnesses are independent measurement instruments of which we know the reliability), it follows that the difference in posterior probability must be due to an increase in coherence.⁵

Furthermore, the example of the defendant repeating his incoherent defense over and over again does not seem to apply to Bovens and Hartmann's measure. For clearly the different testimonies made by the suspect are not independent, and it seems that according to their theory the full testimony by the witness should be considered as one witness report, to be compared with other testimonies from other witnesses.⁶

Next, consider the three measures of coherence as an epistemic virtue. In this case, the intuitive desideratum does seem to apply and it seems we have a genuine counterexample against these measures. But do we? A closer look at the example reveals an interesting point. For in order for the example to be

⁵Of course this argument does not go through anymore if my suggestions for a revision of Bovens and Hartmann's theory of coherence (see chapter 4) are accepted. However, in that case a similar, though less general, argument is still possible. For consider a set consisting of n probabilistically equivalent propositions $S = \{R_1, \dots, R_n\}$ and compare this with a set

$$S_c = \left\{ \bigwedge_{i=1}^n R_i \right\}.$$

In that case the propositions in S and in S_c have the same marginal probabilities, while equation 6.1 still holds.

⁶As Stephan Hartmann has remarked (personal communication).

compelling, it seems that we must assume that the logical relations between the propositions are *known*, so that it is clear to us that one and the same story is told in a number of different ways. But as a general assumption this would certainly be too strong. And absent that assumption neither the intuitive desideratum nor EC can be upheld. The reason for this is that coherence is a hyperintensional notion: how information is presented to us matters to our intuitive judgements with respect to coherence.⁷

But if we accept that coherence is a hyperintensional notion, then it should be totally unsurprising that logically equivalent sets can have different degrees of coherence, at least on any measure of coherence that tries to do justice to our intuitions. Therefore, neither EC nor the intuitive desideratum provide a compelling argument against the measures of coherence as an epistemic virtue.

Nevertheless, the above discussion points to a serious problem for probabilistic approaches to modeling coherence. For a consequence of the Kolmogorov axioms is that if A and B are equivalent, $p(A) = p(B)$. But this can only be an idealizing assumption with respect to our intuitive judgements of probability. This becomes clear if we consider the probability we would assign to two different necessary truths:

1. $A \vee \neg A$;
2. $x^n + y^n = z^n$ has no non-zero integer solutions for x, y and z when $n > 2$,

the second being the famous Fermat's theorem, which was proven in 1993 by Andrew Wiles. Since both are necessary truths, they are equivalent and thus they should have the same (unit) probability. Nonetheless, it seems extremely awkward to argue that everybody who, before the proof in 1993, did not assign both theorems the same probability, was irrational. Clearly, before the proof was completed, it seemed quite reasonable to assign Fermat's theorem a significantly lower probability than the theorem that either A or $\neg A$.

From this a general problem emerges for all probabilistic approaches to the explication of coherence. For such approaches attempt to model coherence with the help of a notion that is intensional only by idealizing assumption. Since in general both our coherence concept and our judgements of probability are hyperintensional, some clashes can be expected between this idealizing assumption and our intuitive judgements of coherence.

⁷It may be instructive to consider an analogy with the concept of explanation. Consider the following example by Douven (2005a) in which two explanations are compared:

1. Lois has nothing to fear, because Clark Kent is Clark Kent.
2. Lois has nothing to fear, because Clark Kent is Superman.

The two explanations of why Lois need not be afraid are equivalent (assuming that 'Superman' is a rigid designator). But, intuitively, only the second explanation really explains. The reason for this is that the concept of explanation is a hyperintensional notion: the way in which information is presented matters for our intuitive judgements with respect to explanation.

This discussion touches on some fundamental issues in analytic philosophy, which cannot receive a full discussion at this point. However, as to the problem of measuring coherence, I believe that there are at least a number of ways in which the problem may be alleviated.

One possible option would be to stipulate that any set whose coherence we could aim to measure should consist of pairwise logically independent propositions. Although initially this may appear to be a rather *ad hoc* move, it seems a quite natural one to make if we seek to analyze a hyperintensional notion in terms of a (merely) intensional notion.

Alternatively, we could simply admit that the measures of coherence have restricted scope. This would not be too bad, either, if we recall that we are after an explication of our intuitive concept of coherence. From this perspective it is certainly not *a priori* the case that there must be *one* formal account of coherence that can handle adequately both sets of propositions that are pairwise logically independent and sets of propositions that are not (all) pairwise logically independent.

Lastly, let me discuss the consequences of EC and the intuitive desideratum for my measure of bootstrap confirmation. As example 6.1 shows, different axiomatizations of a theory may have different degrees of coherence. Moreover, it can readily be seen that accepting a principle similar to EC will have the same type of results as those discussed above. Just consider a theory $T = \{H_1, \dots, H_n\}$ and compare this with a theory $T_c = \{H_c\}$ with

$$H_c \equiv \bigwedge_{i=1}^n H_i. \quad (6.2)$$

Since macho-bootstrapping is not allowed, T_c is not bootstrap confirmed by the evidence even if T is bootstrap confirmed to a very high degree. But whereas coherence is clearly a hyperintensional notion, this is much less clear for confirmation. So how are we to respond to this problem? Well, let me at least give two suggestions for how the problem may be tackled.

Firstly, note 6 in chapter 5 makes the measures \mathcal{B}_m applicable only to theories that are naturally axiomatized. Although no fully satisfactory explication of that notion has yet been given, we can hope that on such an explication the natural axiomatizations of all equivalent theories come out as having the same degree of bootstrap confirmation.⁸ In that case, the problem evidently would no longer arise.

Secondly, we may take seriously the hint at the end of the last chapter and interpret the measures \mathcal{B}_m as measures of the coherence between a theory and the evidence. In that case the above remarks *do* apply and we can at least find

⁸On Gemes's (1993) account of such a notion, whether or not this is the case remains unclear.

an *ad hoc* solution to the problem by making \mathcal{B}_m ($m \in \{\delta, r, c\}$) applicable only to combinations of theoretical and evidential statements in which all sentences are pairwise independent.

6.3 Conclusion

This chapter has discussed two general objections made against the general project of measuring coherence. Although I believe I have shown that the argument by Siebel fails, it is interesting to note that both arguments are concerned with the relation between measuring coherence and deductive logic. However, as this chapter has indicated, such problems could be expected given the fact that while we make the idealizing assumption that probability is not a hyperintensional notion, we would not similarly want to make the same assumption with respect to the concept of coherence.

Moreover, the most awkward consequences can be evaded by simply excluding sets in which not all of the propositions are logically independent from the intended domain of our measures of coherence. Although this does not solve a similar problem for the measure of bootstrap confirmation, it does remove the sting from the arguments with respect to the measures of coherence. Although evidently such a strategy makes these measures less generally applicable, we are left with a large range of cases in which they can still be used to measure the coherence of sets of propositions or theories. One notable drawback of this solution, though, is that some of the counterexamples presented in this thesis no longer go through. For example, the second set in the Tweety examples against Olsson's and Bovens and Hartmann's accounts will be excluded if we accept the above suggestion. Nevertheless, this will not help these measures much: in each of the examples, the logical relations between the propositions were not the reason for the counterintuitiveness of the behavior of the respective measures. Instead, there are similar examples for each of the ones discussed above in which none of the propositions are equivalent to or logically entailed by any of the other propositions. One such alternative is already given in chapter 4, in which I replaced the set $\{B, G, P\}$ by $\{B, F, A\}$.

Chapter 7

Conclusions

In this thesis I have proposed and defended four different explications of our intuitive concept of coherence (five, if we include the measure of bootstrap confirmation). Evidently, it would be more expedient to have only one single measure of coherence. However, in chapter 1 I have argued that this is highly unlikely to occur because the concept of coherence has been used in too many different contexts to allow for a single formal explication, even if we restrict ourselves to the philosophical concept of coherence, as I have done. According to Carnap, for an explication to be as satisfactory as possible, we must be very clear about what the explicandum is. From this it seems to follow that if the explicandum is best characterized by different and contrasting descriptions, we require different explications for each of the different senses that can be distinguished.

Following Carnap's approach to the explication of the concept of confirmation, I have proposed probabilistic measures of coherence as explications of the different senses of the concept of coherence. At times I have availed myself of the more intuitive reasoning employed by Carnap in order to arrive at the most satisfactory explication, at other times I have opted for the desiderata-based strategy that can be found in the works of Kemeny, Oppenheim and Fitelson. Nonetheless, I have not presented a general argument in favor of any of the two approaches. Rather, it appears that when sufficiently many desiderata can be found that are beyond question, the desiderata-based approach provides an objective strategy for arriving at the most satisfactory measure. But if on the other hand none or not enough of such desiderata can be found, it may be more appropriate to resort to some more intuitive considerations.

I have discussed three general types of coherence. All of these have in common that coherence is explicated as a relation between propositions (or between a theory and the evidence) that has three general characteristics: it is a matter of

hanging together, it is a matter of degree and it is weaker than logical entailment but stronger than logical consistency. I have shown that these three characteristics can support a large variety of different explications.

Firstly, there were the explications of coherence as an epistemic virtue. McMullin, for instance, has argued that coherence is one of the virtues that contribute to a theory's explanatory success. In this sense of coherence, the hanging together property has turned out to be pivotal. Indeed, it can be explicated in at least three different ways, which were discussed in chapters 2 and 3. Using these different explications of the hanging together property, these chapters have presented three different explications of this sense of coherence: that of coherence in the sense of mutual support, that of coherence in the sense of relative overlap and that of coherence as an epistemic virtue, where the third was defined as the weighted average of the other two.

Secondly, I have discussed the sense of coherence as a confidence boosting property. According to some versions of the coherentist theory of justification, we can answer the claims of the epistemic skeptic with the help of the concept of coherence. More precisely, if the reports of some witnesses cohere strongly enough, we are entitled to believe these reports, even if the witnesses are highly unreliable individually. In chapter 4, I discussed the impossibility result of Bovens and Hartmann, which shows that coherence cannot be a truth conducive property *per se*, in the sense that a higher coherence implies a higher probability, all else being equal. However, their result does not exclude that coherence is truth conducive in the more restrictive sense of Bovens and Hartmann's version of Bayesian Coherentism. For according to them, the relation of 'being more coherent than' is a quasi-ordering: while in some cases a higher coherence does indeed increase our confidence that the information is true, *ceteris paribus*, in other cases the question must remain unanswered. The vital question in chapter 4 turned out to be how we should interpret the *ceteris paribus* clause. Whereas Bovens and Hartmann argue that the prior probability of the information and the reliability of the witnesses should remain constant, I have shown that if we instead let reliability and the *marginal probabilities* of the propositions remain equal, a very different picture emerges. I argued that the difference function that results from the latter *ceteris paribus* conditions is much more satisfactory, because it does not have the same counterintuitive consequences as Bovens and Hartmann's theory.

Finally, I have attempted to define a quantitative measure of Glymour's theory of bootstrap confirmation. Although this measure is by definition a measure of confirmation rather than of coherence, I argued that it goes at least some way toward explicating the intuitive notion of coherence between a theory and the evidence. More specifically, I discussed Bonjour's claim that foundationalists need a conception of coherence to account for the relation between the foundational and the superstructure beliefs and I proposed that such a conception may be

based on my measure of bootstrap confirmation.

Each of the measures presented in chapters 2–5 has remained tentative in nature. For three of the explications, viz., coherence as relative overlap, as an epistemic virtue and as bootstrap confirmation, I discussed a number of different measures that all seemed to be plausible explications of the intuitive sense of coherence under consideration. I made a tentative case for measures C_d and $\mathcal{E}_{s,o}$ being the most satisfactory explications of coherence as mutual support and as an epistemic virtue, respectively, but I gave no argument in favor of any of the three measures of bootstrap confirmation defined in chapter 5. And although I have not found any measures that are as satisfactory as the measures of coherence as relative overlap defined in chapter 3 and no theories of coherence that are as satisfactory as the difference function presented in chapter 4, it seems quite probable that for these cases, too, many alternatives can be found, some of which may even outperform the measures proposed in this thesis.

Nonetheless, I believe I have convincingly shown that the measures proposed in this thesis constitute more satisfactory explications of our various intuitive concepts of coherence than any of the other measures proposed in the literature. Of the measures proposed by Shogenji and Olsson, I have argued that they fail to take into account all the dependencies between propositions that matter intuitively. For Olsson's measure I have shown how this lack of sensitiveness leads to the Tweety example devised by Bovens and Hartmann and discussed by me in section 3.5. And whereas Fitelson's measure does take into account all of the dependencies that matter intuitively, it fails because it tries to satisfy both the equivalence and the dependence desiderata. This, as I showed in chapter 3, will necessarily lead to the counterintuitive consequence that two sets that seem intuitively almost equally coherent must have a very different degree of coherence. Finally, of Bovens and Hartmann's theory of coherence I have shown that it leaves indeterminate even some intuitively crystal clear cases. And, on top of that, I have shown that their account is actually inconsistent with their definition of the concept of coherence as a confidence boosting property.

Equally importantly, I have shown that none of the counterexamples and general criticisms of measures of coherence presented in the literature so far can undermine the project of measuring coherence or any of the specific explications of coherence proposed in this thesis.

Although these two remarks can in no way diminish the tentative nature of the different explications of coherence proposed, I feel it is safe to conclude that each of the above measures of coherence constitutes a quite satisfactory explication of the respective sense of coherence, or at least a more satisfactory one than any of its contenders.

Quite often in the philosophy of science – on in philosophy generally – the solution of a problem raises more questions than it answers. And this seems

to be the case for the explications of the concept of coherence proposed in this thesis also. One very large topic that has been fully neglected here is the question of truth conduciveness of coherence in any of the other senses of coherence than that of coherence as a confidence boosting property. It must have become clear that none of the other measures of coherence can be considered explications of coherence as a confidence boosting property, for in general it will not be the case that an increase in coherence (as measured by any of these measures) implies an increase in probability (even if we allow for indeterminate cases, as does Bovens and Hartmann's account). But taken literally, to be a property that is conducive to a theory's truth means that it will, in some still to be specified way and under some still to be specified circumstances, increase the probability that a theory is true. And it is not clear that the above explications cannot be truth conducive in this weaker sense of somehow promoting a theory's probability.

In this respect, an interesting result has been derived by Dietrich and Moretti. In their (2005), they show that under some general conditions there are threshold values for the measures of coherence proposed by Olsson and Fitelson such that a degree of coherence above that threshold implies that evidence for one of the hypotheses of that theory is also evidence for any of the other hypotheses. To me it seems that these measures can therefore be called truth conducive in the weaker sense.

One of the problems of Dietrich and Moretti's result is that it only applies to Fitelson's measure for the case of two propositions and, more importantly, that the threshold level will in general be very high.¹ It would therefore be very interesting to find out whether or not stronger results can be derived for the measures presented in this thesis.

A second large topic that I have only briefly discussed is the question of the relation between the concept of coherence as a hyperintensional notion and the concept of probability, which is standardly formalized as an intensional (but not hyperintensional) notion. Many interesting questions remain and it cannot be excluded *a priori* that no other solutions to this problem can be found than the ones I have offered in the previous chapter.

So, doubtless, the last word has not yet been said on the subject of measuring coherence in a mathematically precise way, which is just as it should be for a topic that has emerged only a few years ago.

¹For example, if $p(R) = .6$, $p(E) = .1$ and if $p(R \wedge E) = .1$ (so that E entails R), then the threshold value for Olsson's measure is approximately .94 and for Fitelson's measure approximately .96.

References

- Akiba, K. (2000) 'Shogenji's Probabilistic Measure of Coherence is Incoherent,' *Analysis* 60: 356–359.
- Bonjour, L. (1985) *The Structure of Empirical Knowledge*, Cambridge MA: Harvard University Press.
- Bonjour, L. (1997) 'Haack on Justification and Experience,' *Synthese* 112: 13–23
- Bonjour, L. (1999) 'The Dialectic of Foundationalism and Coherentism,' in J. Greco and E. Sosa (eds.) *The Blackwell Guide to Epistemology*, Oxford: Blackwell, pp. 117–142.
- Bonjour, L. (2002) *Epistemology*, Lanham MD: Rowman and Littlefield.
- Bovens, L. and S. Hartmann (2003a) *Bayesian Epistemology*, Oxford: Oxford University Press.
- Bovens, L. and S. Hartmann (2003b) 'Solving the Riddle of Coherence,' *Mind* 112: 601–633.
- Bovens, L. and S. Hartmann (2005) 'Coherence and the Role of Specificity: A Response to Meijs and Douven,' *Mind*, in press.
- Carnap, R. (1950) *The Logical Foundations of Probability*, Chicago: University of Chicago Press.
- Christensen, D. (1983) 'Glymour on Evidential Relevance,' *Philosophy of Science* 50: 471–481.
- Christensen, D. (1990) 'The Irrelevance of Bootstrapping,' *Philosophy of Science* 57: 644–662.
- Christensen, D. (1997) 'What is Relative Confirmation?,' *Noûs* 31: 370–384.
- Christensen, D. (1999) 'Measuring Confirmation,' *Journal of Philosophy* 96: 437–461.
- Churchland, P. (1985) 'The ontological status of observables: In praise of the superempirical virtues,' in: P. M. Churchland and C. A. Hooker (eds.) *Images of Science*, Chicago: University of Chicago Press, pp. 35–47.
- Dietrich, F. and L. Moretti (2005) 'On Coherent Sets and the Transmission of Confirmation,' manuscript.
- Douven, I. (1999) 'Inference to the Best Explanation Made Coherent,' *Philosophy*

- of *Science* (Proceedings) 66: S424-S435.
- Douven, I. (2002) 'Testing Inference to the Best Explanation,' *Synthese* 130: 355-377.
- Douven, I. (2003) 'On Bayesian Logic,' in: M. C. Galavotti (ed.) *Observation and Experiment in the Natural and Social Sciences*, Dordrecht: Kluwer, pp. 321-326.
- Douven, I. (2005a) 'A Principled Solution to Fitch's Paradox,' *Erkenntnis*, in press.
- Douven, I. (2005b) 'Basic Beliefs, Coherence, and Bootstrap Support,' in: R. van Woudenberg, S. Roeser and R. Rood (eds.), *Basic Belief and Basic Knowledge*, Frankfurt a. M.: Ontos Verlag, pp. 57-76.
- Douven, I. (2005c) 'Empirical Equivalence, Explanatory Force, and the Inference to the Best Theory,' in: R. Festa, A. Aliseda, and J. Peijnenburg (eds.) *Logics of Scientific Cognition: Essays in Debate with Theo Kuipers*, Atlanta GA: Rodopi, in press.
- Douven, I. and W. Meijs (2005a) 'Measuring Coherence,' *Synthese*, in press.
- Douven, I. and W. Meijs (2005b) 'Bootstrap Confirmation made Quantitative,' *Synthese*, in press.
- Dreier, J. (1996) 'Rational Preference: Decision Theory as a Theory of Practical Rationality,' *Theory and Decision* 40: 249-276.
- Duhem, P. (1906/1954) *The Aim and Structure of Physical Theory* (translated by P. Wiener), Princeton NJ: Princeton University Press.
- Earman, J. (ed.) (1983) *Testing Scientific Theories*, Minneapolis: University of Minnesota Press.
- Earman, J. and C. Glymour (1988) 'What Revisions Does Bootstrap Testing Need?,' *Philosophy of Science* 55: 260-264.
- Earman, J. and W. Salmon (1992) 'The Confirmation of Scientific Hypotheses,' in: M. Salmon (ed.) *Introduction to the Philosophy of Science*, Upper Saddle River NJ: Prentice Hall, pp. 42-103.
- Edidin, A. (1983) 'Bootstrapping without Bootstraps,' in Earman (ed.) 1983, pp. 43-54.
- Edidin, A. (1988) 'From Relative Confirmation to Real Confirmation,' *Philosophy of Science* 55: 265-271.
- Eells, E. and B. Fitelson (2000) 'Measuring Confirmation and Evidence,' *Journal of Philosophy* 97: 663-672.
- Eells, E. and B. Fitelson (2002) 'Symmetries and Asymmetries in Evidential Support,' *Philosophical Studies* 107: 129-142.
- Everitt, N. and A. Fisher (1995) *Modern Epistemology - A New Introduction*, New York: McGraw-Hill.
- Fitelson, B. (2001a) *Studies in Bayesian Confirmation Theory*, doctoral dissertation, University of Wisconsin-Madison.
- Fitelson, B. (2001b) 'A Bayesian Account of Independent Evidence with Applica-

- tions,' *Philosophy of Science* 68 (Proceedings): S123-S140.
- Fitelson, B. (2003) 'A Probabilistic Theory of Coherence,' *Analysis* 63: 194-199.
- Fitelson, B. (2004) 'Probabilistic Coherence from a Logical Point of View,' presentation at the London School of Economics, June 2004, see <http://www.fitelson.org/research.htm>.
- Fitelson, B. (2005a) 'Inductive Logic,' in: J. Pfeifer and S. Sarkar (eds), *Philosophy of Science: An Encyclopedia*, New York: Routledge Press, in press.
- Fitelson, B. (2005b) 'Likelihoodism, Bayesianism, and Relational Confirmation,' *Synthese*, in press.
- Fitelson, B. (2005c) 'Bayesian Epistemology [Book Review],' *Mind*, in press.
- Gaifman, H. (1985) 'On Inductive Support and Some Recent Tricks,' *Erkenntnis* 22: 5-21.
- Gemes, K. (1993) 'Hypothetico-Deductivism, Content, and the Natural Axiomatization of Theories,' *Philosophy of Science* 60: 477-487.
- Gemes, K. (1994) 'A New Theory of Content I: Basic Content,' *Journal of Philosophical Logic* 23: 596-620.
- Gemes, K. (1997) 'A New Theory of Content II: Model Theory and Some Alternatives,' *Journal of Philosophical Logic* 26: 449-476.
- Gillies, D. (2000) *Philosophical Theories of Probability*, London: Routledge.
- Glymour, C. (1980a) *Theory and Evidence*, Princeton NJ: Princeton University Press.
- Glymour, C. (1980b) 'Bootstraps and Probabilities,' *Journal of Philosophy* 67: 691-699.
- Glymour, C. (1983) 'Revisions of Bootstrap Testing,' *Philosophy of Science* 50: 626-629.
- Goossens, W. (1976) 'A Critique of Epistemic Utilities,' in: R. Bogdan (ed.) *Local Induction*, Dordrecht: Reidel, pp. 93-113.
- Hacking, I. (2001) *An Introduction to Probability and Inductive Logic*, Cambridge: Cambridge University Press.
- Hage, J. (2004) 'Law and Coherence,' *Ratio Juris* 17: 87-105.
- Halpern, J. (2003) *Reasoning about Uncertainty*, Cambridge MA: MIT Press.
- Hampton, J. (1994) 'The Failure of Expected-Utility Theory as a Theory of Reason,' *Economics and Philosophy* 10: 195-242.
- Hansson, B. (1988) 'Risk Aversion as a Problem of Conjoint Measurement,' in: P. Gärdenfors and N.-E. Sahlin (eds.) *Decision, Probability, and Utility*, Cambridge: Cambridge University Press, pp. 136-158.
- Hempel, C. G. (1965) *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: The Free Press.
- Hetherington, S. C. (1996) *Knowledge Puzzles - An Introduction to Epistemology*, Boulder: Westview Press.
- Horwich, P. (1980) 'The Dispensability of Bootstrap Conditions,' *Journal of Phi-*

- Philosophy* 77: 699–702.
- Howson, C. (2000) *Hume's Problem: Induction and the Justification of Belief*, Oxford: Clarendon Press.
- Howson, C. (2003) 'Bayesian Evidence,' in: M.C. Galavotti (ed.) *Observation and Experiment in the Natural and Social Sciences*, Dordrecht: Kluwer, pp. 301–320.
- Jeffrey, R. (1983) *The Logic of Decision* (2nd ed.), Chicago IL: University of Chicago Press.
- Kaplan, M. (1981a) 'Rational Acceptance,' *Philosophical Studies* 40: 129–145.
- Kaplan, M. (1981b) 'A Bayesian Theory of Rational Acceptance,' *Journal of Philosophy* 78: 305–330.
- Kemeny, J. and P. Oppenheim (1952) 'Degrees of Factual Support,' *Philosophy of Science* 19: 307–324.
- Klein, P. and T. Warfield (1994) 'What Price Coherence?' *Analysis* 54: 129–132.
- Klein, P. and T. Warfield (1996) 'No Help for the Coherentist,' *Analysis* 56: 118–121.
- Kuipers, T. (2000) *From Instrumentalism to Constructive Realism*, Dordrecht: Kluwer.
- Kukla, A. (1998) *Studies in Scientific Realism*, Oxford: Oxford University Press.
- Ladyman, J. (2002) *Understanding Philosophy of Science*, London: Routledge.
- Lehrer, K. (1990) *Theory of Knowledge* (2nd ed.), London: Routledge.
- Levi, I. (1967) *Gambling with Truth*, Cambridge MA: MIT Press.
- Lewis, C.I. (1946) *An Analysis of Knowledge and Valuation*, La Salle: Open Court.
- Maher, P. (1993) *Betting on Theories*, Cambridge: Cambridge University Press.
- McMullin, E. (1984) 'A Case for Scientific Realism,' in: J. Leplin (ed), *Scientific Realism*, Berkeley: University of California Press, pp. 8–40.
- McMullin, E. (1996a) 'Rationality and Paradigm-Change in Science,' in: Horwich, P. (ed), *World Changes: Thomas Kuhn and the Nature of Science*, Cambridge MA: MIT Press, pp. 55–78.
- McMullin, E. (1996b) 'Epistemic Virtue and Theory Appraisal,' in: Douven, I. and L. Horsten, (eds), *Realism in the Sciences: Proceedings of the Ernan McMullin Symposium*, Leuven (Belgium): Leuven University Press, pp. 11–34.
- Meijs, W. (2005a) 'A Corrective to Bovens and Hartmann's Measure of Coherence,' *Philosophical Studies*, in press.
- Meijs, W. (2005b) 'Coherence as Generalized Equivalence,' manuscript.
- Meijs, W. and I. Douven (2005) 'Bovens and Hartmann on Coherence,' *Mind*, in press.
- Merricks, T. (1995) 'On Behalf of the Coherentist,' *Analysis* 55: 306–309.
- Moretti, L. and K. Akiba (2005) 'Probabilistic Measures of Coherence and the Problem of Belief Individuation,' manuscript.
- Niiniluoto, I. (1998) 'Verisimilitude: The Third Period,' *British Journal for the Phi-*

- Philosophy of Science* 49: 1–29.
- Niiniluoto, I. (1999) *Critical Scientific Realism*, Oxford: Clarendon Press.
- Nozick, R. (1981) *Philosophical Explanations*, Cambridge MA: Harvard University Press.
- Olsson, E. J. (1999) ‘‘Cohering With’’ ,’ *Erkenntnis* 50: 273–291.
- Olsson, E. J. (2001) ‘Why Coherence is not truth conducive,’ *Analysis* 61: 236–241.
- Olsson, E. J. (2002) ‘What is the Problem of Coherence and Truth?’ *Journal of Philosophy* 94: 246–272.
- Olsson, E. J. (2005) ‘The Impossibility of Coherence,’ *Erkenntnis*, in press.
- Psillos, S. (1999) *Scientific Realism: How Science Tracks Truth*, London: Routledge.
- Quine, W. V. (1953) ‘Two Dogmas of Empiricism,’ in his *From a Logical Point of View*, Cambridge MA: Harvard University Press, pp. 20–46.
- Rabin, M. (2000) ‘Risk Aversion and Expected-Utility Theory: A Calibration Theorem,’ *Econometrica* 68: 1281–1292.
- Ramsey, F. P. (1926) ‘Truth and Probability,’ in his *The Foundations of Mathematics*, London: Routledge and Kegan Paul, 1931, pp. 156–198.
- Resnik, M. (1987) *Choices*, Minneapolis: University of Minnesota Press.
- Rosenkrantz, R. (1983) ‘Why Glymour is a Bayesian,’ in: Earman, J. (ed), pp. 69–97.
- Satz, D. and J. Ferejohn (1994) ‘Rational Choice and Social Theory,’ *Journal of Philosophy* 91: 71–87.
- Savage, L. (1954) *The Foundations of Statistics*, New York: Wiley.
- Schurz, G. (1991) ‘Relevant Deduction: From Solving Paradoxes Towards a General Theory,’ *Erkenntnis* 35: 391–437.
- Shafer, G. (1976) *A Mathematical Theory of Evidence*, Princeton NJ: Princeton University Press.
- Shogenji, T. (1999) ‘Is Coherence truth conducive?,’ *Analysis* 59: 338–345.
- Siebel, M. (2004) ‘On Fitelson’s Measure of Coherence,’ *Analysis* 64: 189–190.
- Siebel, M. (2005) ‘Probabilistic Measures of Coherence,’ manuscript.
- Spohn, W. (1988) ‘Ordinal Conditional Functions: A Dynamic Theory of Epistemic States,’ in: W. Harper and B. Skyrms (eds.) *Causation in Decision, Belief Change, and Statistics* (Vol. 2), Dordrecht: Reidel, pp. 105–134.
- van Fraassen, B. C. (1980) *The Scientific Image*, Oxford: Clarendon Press.
- van Fraassen, B. C. (1983a) ‘Theory Comparison and Relevant Evidence,’ in Earman (ed.) 1983, pp. 27–42.
- van Fraassen, B. C. (1983b) ‘Glymour on Evidence and Explanation,’ in Earman (ed.) 1983, pp. 165–176.
- Weirich, P. (1986) ‘Expected Utility and Risk,’ *British Journal for the Philosophy of Science* 37: 419–442.
- Weirich, P. (2001) ‘Risk’s Place in Decision Rules,’ *Synthese* 126: 427–441.

- Zadeh, L. A. (1978) 'Fuzzy Sets as a Basis for a Theory of Possibility,' *Fuzzy Sets and Systems* 1: 3-28.

Samenvatting

Het onderwerp van dit proefschrift bevindt zich op het raakvlak van kennisleer en wetenschapsfilosofie. Meer in het bijzonder staat het geven van een precieze analyse van onze intuïtieve notie van coherentie centraal. Er wordt in de wetenschapsfilosofie veel gesproken over de coherentie van een theorie en over de vraag of coherentie in het algemeen een indicator van waarheid kan zijn. Echter, een veel gehoorde klacht in deze discussies is dat het begrip coherentie hopeloos vaag is.

In de laatste jaren zijn er een aantal voorstellen gedaan voor een formele explicatie van het concept coherentie. In het eerste hoofdstuk van dit proefschrift benadruk ik dat dit concept op verschillende intuïties berust en dat we niet uit kunnen sluiten dat elk van deze intuïtieve concepten op een verschillende wijze geformaliseerd dient te worden. Ik onderscheid (zonder te willen beweren dat dit de enige of zelfs de belangrijkste noties van coherentie zijn) drie verschillende betekenissen van het concept coherentie zoals dit in de wetenschapsfilosofie gebruikt is. Deze zijn coherentie als een epistemische deugd, coherentie als een waarheids-bevorderende eigenschap en coherentie als een relatie tussen theorie en evidentie.

Hoofdstukken 2 en 3 geven twee mogelijke explicaties van het concept coherentie als een epistemische deugd, te weten, coherentie als wederzijdse ondersteuning en coherentie als relatieve overlap. In beide gevallen stel ik een probabilistische maat voor en probeer te laten zien dat deze beter aansluit bij onze intuïties met betrekking tot deze vorm van coherentie dan de voorstellen die in de literatuur zijn gedaan.

In hoofdstuk 4 bespreek ik Bovens en Hartmann's explicatie van coherentie als een waarheids-bevorderende eigenschap. Alhoewel het zal blijken dat Bovens en Hartmann's voorstel in tegenspraak is met hun eigen theorie van wat het betekent om een waarheids-bevorderende eigenschap te zijn, zal ik laten zien dat een aantal kleine aanpassingen in hun model tot een nieuwe maat kunnen leiden, die niet tot dezelfde problemen leidt als Bovens en Hartmann's oorspronkelijke voorstel.

In hoofdstuk 5 bespreek ik Glymour's theorie over bootstrap confirmatie. Ik

laat zien hoe de belangrijkste argumenten die tegen bootstrap confirmatie zijn ingebracht, niet langer geldig zijn als we de theorie in probabilistische termen weergeven. Alhoewel het hier een theorie van confirmatie betreft en geen theorie van coherentie, zal ik wel aangeven dat de notie van bootstrap confirmatie in ieder geval aan de basis zou kunnen staan van een explicatie van coherentie als een relatie tussen theorie en evidentie.

In hoofdstuk 6 probeer ik vervolgens een aantal algemene tegenargumenten tegen probabilistische maten van coherentie te ontkrachten.

Curriculum Vitae

Wouter Meijs was born in Nuenen in 1976. He went to school at the Lorentz Lyceum in Eindhoven, which he finished in 1995. From 1995 to 2003 he studied Physics and Philosophy at Utrecht University, while taking additional courses at the Social Sciences Departments of Utrecht University and at the University of Amsterdam. His thesis 'Probabilistic Bootstrap Confirmation' was supervised by Jos Uffink and earned him a degree in Physics (January 2003) and Philosophy of the Exact Sciences (January 2003). At the meantime he worked on a thesis 'International Moral Contract Theories – A Critical Survey,' which was supervised by Ton van den Beld and delivered him a degree in Practical Philosophy (March 2004).

After graduation he started a PhD. research in October 2003 on the basis of his earlier work on Glymour's theory of bootstrap confirmation. He has published in *Mind*, *Philosophical Studies*, *Studies in History and Philosophy of Modern Physics* and *Synthese* as well as in the *Kritisch Denkers Lexicon* and in the edited volume *Denkers van Nu*. He also taught an introductory class on epistemology and refereed for *Studies in History and Philosophy of Modern Physics*, *Synthese*, and *Erkenntnis*.