Let's Reappraise Carnapian Inductive Logic!

Thesis submitted in fulfilment of the requirements of the degree of Doctor of Philosophy in Philosophy

School of European Culture and Languages, University of Kent

Teddy Groves

Total word count: 77523

Acknowledgements

Contents

Chapter 1

Introduction

1.1 Reappraising Carnap

Sometimes the scholarly community comes to realise that it should have a second look at an old idea. Sometimes circumstances change so that the old idea acquires extra salience. Sometimes, for one reason or another, it turns out that an old idea had not previously been understood correctly. When this happens and the old idea is re-examined, prompting a change in the general consensus as to its nature and importance, a reappraisal has taken place.

In recent times scholars of twentieth century philosophy have begun to reconsider their views about the German philosopher Rudolf Carnap and the 'logical positivist' movement that he championed. Carnap and the other logical positivists had previously been understood as continuing the tradition of British empiricists such as Locke, who sought to establish that sensory experience was the foundation of all human knowledge. Lately, works such as Friedman (1999) and Richardson (1998) have argued convincingly that the logical positivists were not primarily interested in this foundational project, and that their work is better understood in the context of Austro-German neo-Kantianism. Logical positivism had been thought of as an apolitical movement, if not one with reactionary tendencies: a well-informed activist once commented "I cant see anything progressive. . . coming out of positivism. It just seems to exclude it completely." Chomsky et al. (1988). Recent scholarship, for example Carus (2010) and Yap (2010), has revealed that many of the logical positivists saw their work as part of an emancipatory social project and that that their ideas can make interesting and useful contributions to current such projects. Close examination of the historical record has also overthrown the previous view that the positivists held academic activities other than physics in disregard. For example, Galison (1990) and Potochnik and Yap (2006) unearth connections between the logical positivists and the Bauhaus movement.

This sustained rethinking of Carnap's place in the intellectual-historical landscape amounts to what might be called a historical reappraisal.

At the same time, many philosophers have begun to argue that some of Carnap's philosophical positions might have greater merit than had previously been supposed. Early examples of this kind of substantive reappraisal can be found in Creath (1991) and Stein (1992), which argue that Carnap might not have been so dogmatic about the controversial topics of analyticity and reductionism as had previously been believed. There has been renewed interest in Carnap's method of explication, expressed by Maher (2007) and Olsson (2014). Peregrin (2015), Awodey (2015) and Restall (2002) have reappraised Carnap's contributions to the philosophy of logic. Psillos (2000) reappraises Carnap's views on formalising scientific theories. Awodey and Carus (2001) reappraise Carnap's views on the philosophy of mathematics. Lutz (2012), Leitgeb (2013) and Carus (2014) reappraise Carnap's meta-philosophical stance.

It seems that most, if not almost all, of Carnap's philosophical positions are currently being historically and substantively reappraised. The most striking exception to this general trend is the topic that most occupied Carnap towards the end of his career and is the main subject of this thesis, namely Carnapian inductive logic.

From the 1940s until his death in 1970, Carnap and several co-researchers conceived and developed a distinctive method of formalising inductive reasoning, which they claimed was strongly analogous to deductive logic. While this research has been discussed in the recent philosophical literature[1], and some aspects of it have been extended, principally in the work of a group of inductive logicians based in Manchester, culminating in Paris and Vencovská (2015), it has not yet been reappraised. Among philosophers, the prevailing consensus remains substantially the same as it was when Carnap's publications on inductive logic were first criticised in works such as Goodman (1946), Toulmin (1953), Putnam (1963), Salmon (1967) and Lakatos (1968a).

This thesis makes the case that Carnapian inductive should not remain the exception to the general trend towards reappraising Carnap's philosophy. It argues that the time is ripe for both a historical and a substantive reappraisal and begins to carry out both of these tasks.

I aim to begin a historical reappraisal by proposing and defending a new account of Carnapian inductive logic. This account challenges the current understanding of both the nature of Carnapian inductive logic and its relationships with other aspects of Carnap's philosophy and the work of other philosophers.

On a substantive level, this thesis argues that Carnapian inductive logic is a viable research programme that has much to offer contemporary philosophical debates. In what follows I argue that its most prominent critiques can be

---

[1]See, for example, Skyrms (1996), Zabell (2011) and Fitelson (2005), Glaister (2007), Hawthorne (2014), Uebel (2012), Wagner (2011), Romeijn (2005) and Romeijn (2011).

overcome, that it compares favourably with rival research programmes and that it can usefully be applied to a modern debate in the philosophy of statistics.

1.2 Structure

The structure of this thesis is as follows. In chapter 2, I explain the key features of Carnapian inductive logic and describe its development. I construe Carnapian inductive logic as a scientific research programme in Lakatos's sense of the term and argue that it was progressive according to Lakatos's methodology. Next, in chapter 3, I examine the principal critiques of Carnapian inductive logic and argue that each of these can be overcome. Chapters 4 and 5 compare Carnapian inductive logic with related research programmes arising from subjective and objective Bayesian epistemology, arguing that it has advantages over both. In chapter 6 I present an application of Carnapian inductive logic to a contemporary philosophical debate in the philosophy of statistics. Chapter 7 concludes, summing up my historical and substantive claims and setting out a future research agenda.

The following subsections summarise the main points of each chapter.

Chapter 2: Carnapian inductive logic

This chapter presents my main historical claims about Carnapian inductive logic and evaluates it according to Lakatos's methodology of scientific research programmes. It shows how Carnapian inductive logic can be thought of as a scientific research programme, describes how it changed through the three decades during which Carnap worked on it, and argues that it was 'progressive' according to Lakatos's methodology.

I begin by briefly describing the development of Carnapian inductive logic and some related research programmes. Next I introduce Lakatos's methodology and explain how, with minor modifications, it can be applied to Carnapian inductive logic.

I then introduce the key features of Carnapian inductive logic, characterising them as the research programme's 'hard-core' commitments:

1. **Explication** The first commitment is to the overall goal of formalising, or in Carnap's terms 'explicating', inductive assumptions in a way that is useful in everyday life and science.

2. **Systems of inductive logic** The second commitment is to using a certain kind of formal tool to explicate inductive assumptions, namely 'systems of inductive logic'.

3. **Tolerance** The third commitment is to tolerance, at the philosophical level. According to Carnapian inductive logic, the question of whether a system of inductive logic should be adopted is fundamentally practical, not philosophical.

I also classify the other parts of Carnapian inductive logic according to Lakatos's methodology and document the main ways in which the research programme changed as it developed.

This is not the first time that Carnapian inductive logic has been analysed using the methodology of scientific research programmes: Lakatos (1967a) and Niiniluoto (2011) have also done so. However, my conclusions are different because they are based on a novel reading of Carnapian inductive logic; the key difference is my claim that it was committed to tolerance. I defend various aspects of this reading in an appendix. The way I present Carnap's technical apparatus draws considerably from the presentation in Paris and Vencovská (2015), though it is also novel in some ways, and preserves most of Carnap's original terminology.

Based on my reading, I argue that Carnapian inductive logic had the features that, according to Lakatos, characterise 'progressive' research programmes. It steadily expanded its stock of auxiliary claims, broadly stayed within the parameters it originally set for itself and had the potential to be useful in everyday life and science.

I conclude that, while it is interesting that Carnapian inductive logic was progressive in this sense, much work remains to be done. Both a historical and substantive reappraisal require engagement with the many critiques of Carnapian inductive logic. In addition, in order to make a convincing case that Carnapian inductive logic has something to offer present-day philosophy, it must be established how it compares with rival research programmes, and if it has any useful contemporary applications. The remaining chapters seek to fill in these gaps.

Chapter 3: Critiques

This chapter addresses what I consider to be the most important critiques of Carnapian inductive logic, arguing that none of them succeeded in showing that it was unviable.

Lakatos
Lakatos claimed, contrary to my analysis in chapter 2, that Carnapian inductive logic was a degenerate, rather than a progressive, research programme. Lakatos claimed that Carnapian inductive logic abandoned a hard-core commitment, thereby committing what the methodology of scientific research programmes considered an unforgiveable sin. I argue that Lakatos made a crucial mistake in his reading of Carnap's earlier work, leading him to mistakenly attribute

to early Carnapian inductive logic the goal of finding true axioms. Whereas Lakatos's critique depends on Carnapian inductive logic having abandoned this goal, I show that Carnapian inductive logic never had it in the first place.

Grue

Next, I consider the 'grue' problem. I present two versions of this prominent objection: the first claims that Carnapian inductive logic cannot formalise inductive assumptions about 'projectability', whereas the second claims that Carnapian inductive logic cannot show that any one such assumption is more justified than any other. I argue, contrary to the first kind of objection, that Carnapian inductive logic can formalise distinctions between projectability and non-projectability. The second objection, I contend, sets the bar too high. Neither Carnapian inductive logic nor any other intellectual endeavour has yet shown how to justify assumptions about projectability, a goal to which Carnapian inductive logic never aspired.

Logical independence of distinct events

The third critique that I consider is by Salmon, who argued that Carnapian inductive logic cannot simultaneously formalise two intuitive inductive assumptions. These are the assumption of positive instantial relevance, according to which instances of experimental results make future results of the same kind more plausible, and an assumption that Salmon called 'Hume's principle', according to which all distinct events are logically independent. I present an inductive logical axiom that, I claim, formalises Hume's principle. As Salmon postulated, it is inconsistent with the axiom that is usually chosen to formalise positive instantial relevance. I highlight critical responses to Salmon's essay by two of Carnap's co-workers, Richard Jeffrey and John Kemeny, that I think have received insufficient attention in the literature on Carnapian inductive logic. I agree with Kemeny's assessment that Carnapian inductive logic need not formalise both Hume's principle and instantial relevance at the same time.

Other critiques

I then address in less detail some critiques that I think are less important. Contrary to a criticism made by Popper, I argue that Carnapian inductive logic need not formalise inductive assumptions about universal laws in order to be successful. Carnapian inductive logic has been criticised for formalising assumptions of exchangeability, according to which the order of observations carries no information: these are seen as related to the controversial principle of indifference. I argue that, whatever their philosophical status, Carnapian inductive logic can legitimately formalise exchangeability assumptions because they are ubiquitous in statistics. Finally, I address a claim by Putnam that Carnapian inductive logic is inferior with respect to long-run reliability to an alternative formal method that he devised. I argue, following Kelly, Juhl and Glymour, that this is not the case.

I conclude that Carnapian inductive logic survives its most prominent critiques.

Chapter 4: Subjective Bayesian inductive logic

The next stage of my thesis considers two similar research programmes, arguing that Carnapian inductive logic has some advantages over both.

The first is subjective Bayesian inductive logic, a programme that has recently been defended by Colin Howson in a series of works. I construe subjective Bayesian inductive logic as similar to Carnapian inductive logic, except that it rejects the latter programme's commitment to tolerance, while embracing alternative hard-core commitments inspired by subjective Bayesian epistemology. These take the form of the following claims about rational states of partial belief:

- **Probabilistic necessity** Every rational state of partial belief can be represented by a probability function.

- **Probabilistic sufficiency** Every state of partial belief that can be represented by a probability function is rational.

- **Conditionalisation** The way rational agents take evidence into account can be represented using the 'conditional probability' method.

I discuss in detail the epistemological motivation for probabilistic necessity and sufficiency, arguing that neither of these claims is conclusively established by currently available arguments. I argue that subjective Bayesian inductive logic's commitment to conditionalisation is methodologically disadvantageous.

I consider three kinds of argument for probabilistic necessity: betting arguments, axiomatic arguments and accuracy arguments. I claim that betting arguments are thrown into doubt by objections arising from Weatherson (2003) and Schick (1986), while axiomatic arguments have struggled to overcome an objection first highlighted by Halpern (1999a). I claim, contrary to all accuracy arguments, that rational agents may sometimes knowingly choose an inaccurate state of belief. For example, they may do so in order to benefit their epistemic community. In addition, I argue against some properties that have been proposed as legitimacy conditions for measures of inaccuracy. While some of these arguments have appeared elsewhere, I believe that others, including my objections to the putative legitimacy conditions continuity and proposition-neutrality, are novel.

Before concluding my discussion of probabilistic necessity, I consider whether, considered togther, the different kinds of argument for this claim combine into a compelling collective argument. I argue that they are not, as the many similarities between the three main kinds of argument makes the fact that they share the same conclusion unsurprising.

I also consider three kinds of argument for probabilistic sufficiency: grue-based arguments, accuracy-dominance arguments and a logic-based argument due to Howson. Grue-based arguments assert that genuine rationality constraints must

be interpretation-neutral: the axioms that formalise them must always formalise rationality principles, regardless of any interpretative considerations. I claim, contrary to this kind of argument, that there are genuine interpretation-sensitive rationality constraints. Accuracy-dominance arguments claim that all states of partial belief that are not dominated with respect to their inaccuracy are rational, and use this claim to build an argument. I dispute the claim, arguing that irrational states of belief can be un-dominated. Finally, Howson's logic-based argument defends probabilistic sufficiency on the grounds that axioms that are stronger than the probability axioms fail to satisfy criteria that, he claims, are required for genuine logicality. I present a stronger set of axioms that, I claim, either satisfy Howson's criteria or show that they are unsound.

Finally, I consider subjective Bayesian inductive logic's commitment to conditionalisation. Whereas Carnapian inductive logic can formalise all inductive assumptions that can be captured by subjective Bayesian inductive logic, conditionalisation prevents the reverse from being true. As a result, Carnapian inductive logic is the more flexible research programme.

In light of these arguments, I compare Carnapian inductive logic with subjective Bayesian inductive logic. Carnapian inductive logic, which was not committed to either probabilistic necessity or probabilistic sufficiency, has the advantage of being less philosophically controversial than subjective Bayesian inductive logic. From a methodological point of view, Carnapian inductive logic is more flexible, whereas subjective Bayesian inductive logic is simpler in a certain sense.

I conclude that the balance of advantages favours working within Carnapian inductive logic rather than subjective Bayesian inductive logic for the majority of applications.

Chapter 5: Objective Bayesian epistemology and inductive logic

In this chapter I investigate how objective Bayesian epistemology can be connected with inductive logic. Focusing on the form of objective Bayesian epistemology defended by Jon Williamson, I consider two options. The first option is to construct a distinctive objective Bayesian inductive logic, while the second is to formalise the norms of objective Bayesian epistemology within Carnapian inductive logic.

The objective Bayesian inductive logic that I consider is similar to Carnapian inductive logic, except that it is committed to the norms of objective Bayesian epistemology, as outlined in Williamson (2010). Since these include probabilistic necessity, I argue that objective Bayesian inductive logic shares a philosophical disadvantage, compared to Carnapian inductive logic, with subjective Bayesian inductive logic. In addition, I argue that objective Bayesian inductive logic has a methodological disadvantage, as it cannot easily represent certain kinds of evidence. I illustrate this problem with an example involving evidence according to which one kind of observation is positively relevant to another. On the

other hand, objective Bayesian inductive logic has an important methodological advantage compared to Carnapian inductive logic, as it affords less discretion to the inductive logician. I conclude that objective Bayesian and Carnapian inductive logic can each be useful, depending on the application.

In order to formalise the norms of objective Bayesian epistemology within Carnapian inductive logic, the latter discipline must be enriched. Following a suggestion in Skyrms (1985), I sketch how this could be done in principle, showing how objective Bayesian epistemology can be thought of as defining a Carnapian system of inductive logic. I argue that carrying out this 'assimilative approach' would be worthwhile as it would both clarify the debate about the status of the norms and allow pragmatic arguments, as opposed to epistemological ones, to be given for and against them. I note that, while the assimilative approach would be somewhat disconnected from the rest of Carnapian inductive logic, this issue does not fatally undermine it.

I conclude by claiming that both approaches to relating objective Bayesian epistemology and inductive logic are viable and perhaps even complementary.

Chapter 6: Applications of Carnapian inductive logic

The final part of my argument considers ways in which Carnapian inductive logic can be applied in order to solve contemporary philosophical problems. My main case study is an argument in the philosophy of statistics. I argue that, in general, it is helpful to conceive of statistical model-choice as Carnap conceived of choices between systems of inductive logic. Doing so allows this activity to be described in a way that does justice to the actual practice of statistical research while at the same time being philosophically well-grounded. In order to make this point I discuss in detail a recent argument from Gelman and Shalizi (2012), which claims that subjective Bayesian inductive logic cannot capture the kind of reasoning that practising statisticians employ when choosing between statistical models. Gelman and Shalizi go on to claim that thinking about such choices within a Popperian philosophical framework would be preferable. I agree with the first claim, but not the second: the Popperian framework is neither philosophically well-grounded nor capable of doing justice to statistical practice. In contrast, Carnapian inductive logic does much better.

Chapter 7: Conclusion

I conclude by summing up the present state of the reappraisal of Carnapian inductive logic. I claim that new historical and substantive positions have emerged, and that there is a strong case for continuing to reappraise Carnapian inductive logic, a research programme which succeeds according to Lakatos's criteria, survives its critiques, has advantages compared to its rivals and has interesting applications to contemporary philosophical debates. I end by outlining

a research agenda according to which this reappraisal may be completed.

1.3 Limitations of this analysis

I believe that this thesis presents a compelling case for reappraising Carnapian inductive logic, and makes useful progress towards such a reappraisal. However, it is limited in several ways: I list a few of the most conspicuous limitations below.

Sources

I have not been able to take into account material that Carnap wrote in his native German but has not yet been translated. Fortunately, most of Carnap's published work on inductive logic was written in English: see section 2.2 in the next chapter for a summary of the published works that I have drawn on. One important exception is Carnap and Stegmüller (1959), which was written in German and, to my knowledge, has not yet been translated into English. Being able to take this work into account would doubtless have improved the presentation given here.

It would also have been useful for the sake of context to have been able to analyse Carnap's unpublished writing in German on the topic, as well as earlier German-language publications.

Secondly, I was unable to take advantage of many handwritten notes by Carnap on inductive logic that are available at the Pittsburgh archives. Carnap wrote many of these using an unusual shorthand method—the now very unusual 'Stolze-Schrey' method—to which he reportedly added several personal modifications.

Philosophical questions about logic

Some of the philosophical controversy that has surrounded Carnapian inductive logic since its inception has had to do with its use of the term 'logic'. Various authors have been dubious about the very idea of inductive logic on the grounds that inductive reasoning falls outside of the proper purview of logic.

In this regard I devote a section of the appendix to chapter 1 to Carnap's use of the term 'logic'. I argue that Carnap had principled reasons for using the term, based on his general conception of logic and analogies that he drew between inductive and deductive logic.

However, I do not assess whether Carnapian inductive logic should properly be regarded as part of logic according to other conceptions, or engage at length with philosophical discussion of which conception is best.

Although situating inductive logic within these philosophical debates is an interesting task for future research, I felt that that it would not improve this thesis's dialectic. Ultimately, whether or not Carnapian inductive logic is part of logic does not determine whether it is worth reappraising. The arguments I present do not appeal to any special status of logicality: hopefully even readers who are not convinced that Carnapian inductive logic is part of logic will be persuaded despite their terminological reservations.

Critique by Popper

While I address a criticism by Popper of Carnapian inductive logic according to which it does not correctly formalise reasoning involving general laws, I leave many of his other critical remarks undiscussed. There are many such remarks: for example, (Popper, 1959, p. 409), Popper claimed that Carnapian inductive logic was totally unviable, apparently for reasons unrelated to its treatment of general laws. I have omitted discussion of this criticism because I felt that Popper's point remained unclear. In addition, Popper had a separate extended debate with Carnap about whether the term 'confirmation' should be understood in an absolute or relative sense. I avoided discussing this debate because I felt that, whatever the outcome, the question of Carnapian inductive logic's reappraisability would not be greatly affected. The reader is directed to Michalos (1971) for a full account of the Popper-Carnap controversy.

Non-Lakatosian evaluative frameworks

In Chapter 2 I claim that Carnapian inductive logic does well according to Lakatos's methodology of scientific reseach programmes, arguing that this shows that a reappraisal might be in order. While I note that Lakatos's evaluative methodology is controversial and has prominent rivals, I do not address these controversies at length or investigate how well Carnapian inductive logic does according to any rival methodologies. Instead I offer the reader additional, independent reasons to believe that Carnapian inductive logic merits substantive reappraisal, which I hope will outweigh any doubts that they have about Lakatos's methodology.

For the sake of completeness, and in order to come to a definitive view as to whether Carnapian inductive logic was a successful research programme, it would be helpful for future research to assess how it ought to be assessed and how well it does according to alternative, non-Lakatosian evaluative frameworks.

Chapter 2

Carnapian inductive logic

2.1 Introduction

If recent analytic philosophy has one settled opinion about Carnapian inductive logic, it is that it was not successful. Even before the publication of its first comprehensive statement in Carnap (1950b), Carnapian inductive logic had already been influentially critiqued by Goodman in Goodman (1946) and Goodman (1947). As Carnapian inductive logic developed during the 1950s and 1960s, distinguished philosophers continued to denounce it. A general idea of the reception of Carnapian inductive logic in the philosophical community can be gleaned from Popper (1968), Toulmin (1953), Putnam (1963), Salmon (1967) and Lakatos (1968a).

Since Carnap's death in 1970, the philosophical community has continued to produce critiques of his project, some of which we will encounter in chapter 3.

Support for Carnapian inductive logic has been less emphatic. Carnap's co-workers often defended the programme as it progressed, or in its immediate aftermath, but subsequently either went on to work on different projects like Kemeny and Gaifman or took substantially different points of view like Jeffrey and the Finnish school of inductive logicians. While there have been sympathetic contemporary accounts, for example by Zabell (2011), Maher (2010) and the Manchester school, these have tended not to challenge the critical ones directly.

At least within philosophy, then, there is a consensus that Carnapian inductive logic was unsuccessful. But what is, or was Carnapian inductive logic? Neither the critical nor the supportive literature is always entirely clear about this question. Even Carnap's own writings on the subject are somewhat inaccessible, thanks to their length and Carnap's idiosyncratic presentational and

stylistic choices. This chapter aims to improve on the status quo by presenting a comprehensive and accessible account of Carnapian inductive logic.

In what follows I defend a novel assessment of the nature and aims of Carnapian inductive logic. I claim that Carnapian inductive logic is best construed as a research programme, that is, as a movement in the history of science. I attempt a taxonomy of Carnapian inductive logic according to Lakatos's methodology of scientific research programmes. I identify hardcore commitments that characterised Carnapian inductive logic, less important secondary commitments that could be modified or discarded if necessary and 'heuristics' that guided researchers, as well as describing how these features changed as the research programme developed.

Based on this account, I argue that Carnapian inductive logic was a progressive research programme according to Lakatos's criteria.

Before beginning to apply Lakatos's methodology, I will briefly describe the history of Carnapian inductive logic and some related research programmes.

2.2 Historical sketch

Along with several co-workers, Rudolf Carnap studied inductive logic between around 1940 and 1970.

This research is largely encapsulated in three major works. The book *Logical Foundations of Probability*, Carnap (1950b), was aimed at philosophers. It explains what Carnap saw as the overall task of inductive logic, tentatively proposes a particular inductive logical system and attempts to answer some questions about its philosophical and mathematical status. The monograph *The Continuum of Inductive Methods*, Carnap (1952a), attempted to explain a slightly modified form of this system to practising scientists and apply it to some problems in statistics. Finally, *A Basic System of Inductive Logic*, published in two parts after Carnap's death, as Carnap (1971a) and Carnap (1980), documents the programme's later phase.

The intermediate stages of Carnapian inductive logic are discussed in Carnap and Stegmüller (1959), and in Carnap (1963b).

In addition, Carnap published many journal articles on the topic of inductive logic. Aside from Carnap (1945a), which first introduced his programme, these articles largely supplement the longer works. They include Carnap (1945b), Carnap (1946), Carnap (1947b), Carnap (1947c), Carnap (1948), Carnap (1951), Carnap (1953), Carnap (1966a) and Carnap (1968).

Carnap and his co-workers discussed inductive logic in some essays about on other topics, such as the use of 'meaning postulates' in deductive logic Carnap (1952b), the measurement of information Bar-Hillel and Carnap (1953) and

statistical mechanics Carnap and Shimony (1977). Finally, a substantial body of unpublished material on inductive logic by Carnap—correspondence, notes and manuscripts—survives, some of which I cite below.

After Carnap's death in 1970, many of his co-workers, including Yehoushua Bar-Hillel, Haim Gaifman and John Kemeny, wrote mostly about different topics. Others, such as Richard Jeffrey and Abner Shimony, continued to write about the topics that they had focused on as Carnapian inductive logicians, but took a somewhat different perspective.

2.2.1   Related programmes

While Carnapian inductive logic can, I believe, be analysed based on the written material referred to above, it is also important to mention some related research programmes that it either influenced or was influenced by.

Cambridge authors
Though, as we shall see later, the relationship between Carnap's views and those of Keynes is not entirely straightforward, Carnap drew inspiration for his work on inductive logic from Keynes (1921). Carnap was also influenced by the writing on inductive reasoning of several other Cambridge authors, namely Wittgenstein, Harold Jeffreys and Frank Ramsey. Ironically, another Cambridge author, W.E. Johnson, anticipated much of Carnap's work, though Carnap does not seem to have read Johnson's work until well after starting his own research programme.

The Finnish school
Carnap's programme was both pre and post dated by a related programme which Niiniluoto calls "the Finnish school of induction". Starting with the work of Eino Kaila, several Finnish authors including von Wright, Hintikka and Niiniluoto pursued research into inductive logic.

This research programme and its relationship with Carnapian inductive logic are documented in Niiniluoto (2011). Niiniluoto notes that, at least on a sociological level, this tradition was to a large extent distinct from Carnapian inductive logic, having plotted an independent course from the work of Keynes.

The Manchester school
Since the mid 1980s a group of scholars based in Manchester has carried out research into the mathematical structures that were the main subject matter of Carnapian inductive logic, and their generalisations. This research has extended the scope of Carnapian inductive logic, revealed new facts about alternative systems and presented Carnap's work in a much more accessible and modern form. Much of this research is documented in Paris and Vencovská (2015).

2.3 The methodology of scientific research programmes

In order to begin a historical reappraisal of Carnapian inductive logic, it might be best to discuss it from a neutral perspective, identifying important claims, influences and changes without commenting on their merit. However, this kind of analysis would not be sufficient for our second purposes of substantive reappraisal. In order to determine whether Carnapian inductive logic still has something to offer, we need a philosophical framework within which it is possible to distinguish good, reappraisal-worthy research programmes from bad ones that do not deserve renewed attention.

I submit that Lakatos's methodology of scientific research programmes, as articulated in Lakatos (1968b) and Lakatos (1970), is a sensible first candidate for such a framework. Lakatos's methodology sets out a taxonomy according to which scientific research programmes can be classified, describes what they do and sets out evaluative criteria according to which some research programmes are good, or 'progressive', whereas others are bad and 'degenerate'.

In this section I outline the main features of Lakatos's methodology and argue that it is a good starting point for an assessment of Carnapian inductive logic.

2.3.1 Lakatos's methodology

The methodology of scientific research programmes analyses scientific research programmes in terms of two kinds of rule: the 'negative heuristic' and the 'positive heuristic'.

A research programme's negative heuristic identifies a 'hard core' of dogmas that constitute the research programme's central commitments and stipulates that these should not be revised under any circumstances. Only secondary commitments, those identified by the negative heuristic as part of the programme's 'protective belt', should ever be changed. The negative heuristic should provide researchers with instructions as to how they should react when confronted with uncooperative facts that might tempt them to drop one of their hard core commitments.

A research programme's positive heuristic stipulates in advance a strategy for devising and modifying secondary commitments so as to account for more and more facts and eventually make novel predictions.

Lakatos saw Newtonian gravity as an apt example of a scientific research programme. Its hard core consisted of Newton's laws of motion, together with certain other unarticulated commitments[1]: it's negative heuristic prohibited these from being modified. Newtonian gravity's protective belt consisted of a

---

[1]See (Lakatos, 1968b, p.172).

collection of secondary commitments which, together with the hard core dogmas, determined a collection of models of planetary movements. Its positive heuristic gave Newtonian researchers instructions as to how to improve these models, altering and adding secondary commitments along the way, so that they could eventually be used to make novel predictions without contradicting any core Newtonian dogmas.

Besides this taxonomy of a research programme's components, Lakatos's methodology also puts forward standards of success and failure. A research programme is successful, or 'progressive' if it shows 'theoretical progress', 'empirical progress' and 'heuristic power'. Theoretical progress consists in making more and more secondary commitments. Empirical progress occurs if, at least intermittently, discoveries are made that would not have been possible without these new commitments. Finally, in order for a research programme to have heuristic power, changes in its protective belt must not be 'ad-hoc'. Lakatos thought that a change to a research programme could be ad-hoc in three specific ways[2]: by failing to add genuine empirical content, by retrospectively accounting for the discoveries of rival research programmes or by not being in the spirit of the programme's positive heuristic.

If a research programme fails to satisfy any of these criteria, then it is 'degenerate', according to the methodology of scientific research programmes.

See (Lakatos, 1968b, § 2(c) and § 3) and (Lakatos, 1968a, § 3) for Lakatos's discussions of theoretical and empirical progress. See (Lakatos, 1968b, p.174) for comments on heuristic power and (Lakatos, 1970, p.125, footnote 34) for a more detailed discussion of what it means for changes to be ad-hoc.

Newtonian gravity was progressive, according to Lakatos, because researchers following its heuristics managed to build steadily more complicated models of planetary movements, which they eventually used to make novel discoveries about the planets. Moreover, according to Lakatos, the way in which Newtonian researchers improved their models was anticipated by the pattern of Newton's initial research, and was therefore in the spirit of Newtonian gravity's positive heuristic.

2.3.2 The question of truth-directedness

There are two potential issues which might make the reader wary of my plan to apply Lakatos's methodology of scientific research programmes to Carnapian inductive logic.

First, the reader might wonder why it is necessary to apply Lakatos's methodology to Carnapian inductive logic, given that Lakatos himself already did so in Lakatos (1968a), coming to the conclusion in Lakatos (1973) that Carnapian inductive logic showed "all the characteristics of a degenerating programme".

---

[2]See (Lakatos, 1970, p.125, footnote 34).

Unfortunately, we shall see in section 3.2 that Lakatos made some crucial errors in his analysis of Carnapian inductive logic which invalidate this conclusion - as a result further investigation is called for.

A second problem has to do with truth-directedness, and must be addressed before going further. Lakatos's methodology stresses the importance of truth-directed propositions, that is, propositions that aim to make correct claims about the nature of the world. Lakatos characterised a research programme's hard-core as consisting of such claims. For example, he characterised the hard core of Newtonian gravity as consisting of Newton's laws of motion together with his law of gravitation[3] and equated the hard core of algebraic topology with that programme's formal axioms[4]. In addition, Lakatos defined empirical progress so that whether or not a research programme is progressive depends on whether its truth-directed predictions are verified.

This presents a problem for our analysis, as Carnapian inductive logic's commitments and predictions were not straightforwardly truth-directed. Rather than aiming to describe the world correctly, I shall argue, Carnapian inductive logic aimed to be useful in everyday life and science. Rather than making commitments to the truth of particular claims about the world, Carnapian inductive logic made commitments to the usefulness of particular methodological prescriptions. For this reason Lakatos's methodology does not apply to Carnapian inductive logic if it is taken overly literally.

Lakatos himself was aware of this kind of obstacle to the application of his methodology to Carnapian inductive logic. Nonetheless, he thought that the obstacles could be surmounted, and that Carnapian inductive logic could be shown to be a degenerate research programme. Lakatos sought to solve the problem in an unpublished essay entitled 'The Criticism of Carnap's Research Programmes' by identifying Carnapian inductive logic as an 'a priori' kind of research programme, in contrast to the customary 'empirical' kind.

> Carnap's research programme is not a scientific (or empirical) research programme: it does not attempt to set up theories that can clash with facts. Carnap's inductive logic aims not at an empirical theory but at an a priori normative theory.
> (Lakatos, 1967a, p.12)

According to Lakatos's reading of Carnapian inductive logic, which is also apparent in Lakatos (1968a), it did not aim to find 'empirical' truths that can in principle be verified, but rather to find 'a priori' truths that can only be accessed by means of intuition. We shall see in the discussion of Lakatos (1968a) in the next chapter that Lakatos was mistaken to describe Carnapian inductive logic in this way, since it was not even truth-directed in an 'a priori' way.

A better way to unite Carnapian inductive logic and the methodology of scien-

---

[3]See (Lakatos, 1968b, p. 169).
[4]See (Lakatos, 1978b, p.96).

tific research programmes is to make the latter liberal enough to accommodate research programmes with non-truth-directed hard core commitments. Corfield argues that precisely this approach would improve Lakatos's methodology:

> A Newtonian if asked whether he thought the law of gravitation and Newtons three Laws of motion were true, might well have answered affirmatively, as would a seventeenth-century mathematician have responded to a similar question concerning the axioms of Euclidean geometry. But were we to ask a mathematician today whether he considered a system of axioms to be true, he would most probably tell us that he thought it had interesting models or that it described an important construction. Perhaps the distance between the modern mathematician and scientist in this respect is not so great in that the physicist is likely nowadays to say that she thinks her theory a good model of some aspect of the universe. Thus, Lakatos ought to have allowed a wider notion of hard core to allow research programmes any chance of success in accurately portraying science.
> (Corfield, 2003, p. 182)

Corfield claims that allowing the inclusion of "higher-level beliefs and aims" in research programmes' hard cores is necessary in order for Lakatos's methodology to describe mathematical and scientific research adequately. Following this approach, I will apply a slightly modified version of Lakatos's methodology, allowing hard core commitments not to be truth-directed. This divergence, I claim, is minor enough that I can still be said to be applying the methodology of scientific research programmes.

2.4 Hard core commitments

The most important step towards analysing Carnapian inductive logic according to Lakatos's methodology of scientific research programmes is to identify its hard-core commitments. In this section I identify three main methodological prescriptions which, I argue, play this role. These include commitments to:

- **Explication** Producing useful formal analogues or 'explications' of inductive assumptions

- **Systems of inductive logic** Using a certain kind of formal tool, namely 'systems of inductive logic'

- **Tolerance** Taking a philosophically tolerant attitude to the task of evaluating whether or not attempts at explication have been successful.

I shall now discuss these commitments, and their textual bases, in more detail.

2.4.1   Explication

The main goal of Carnapian inductive logic was to produce useful formal replacements for assumptions about inductive reasoning.

Inductive assumptions assert rules as to how uncertain reasoning ought to proceed. An inductive assumption might assert, to take an important and frequently adopted example, that instances of a certain kind of observation should render future instances of the same kind more plausible. A formal replacement for an inductive assumption restates it, or at least approximates it, in a formal language.

The fact that producing formal replacements for inductive assumptions was the main goal of Carnapian inductive logic throughout its development is a key claim of my thesis. That it was an aim of late Carnapian inductive logic can clearly be discerned from the following quotation:

> Now it is just the purpose of IL to give explicit rules for inductive reasoning, so that a subject can, whenever he regards it worth the effort, choose to apply a rule-directed procedure instead of the common instinct-directed procedure.
> Carnap (1967b)

The first six sections of Carnap (1950b), the main statement of early Carnapian inductive logic, are devoted to explaining Carnap's goal of producing useful formal replacements for inductive assumptions. In these sections Carnap provides a general philosophical account of the process of producing useful formalisations, which he called 'explication'. Carnap describes different kinds of explication, proposes criteria for evaluating formal replacements or 'explicata' and argues that explication has played an important role in the history of ideas.

The general nature of this discussion can be gleaned from the following quotation:

> The task of *explication* consists in transforming a given more or less inexact concept into an exact one or, rather, in replacing the first by the second. We call the given concept (or the term used for it) the *explicandum*, and the exact concept proposed to take the place of the first (or the term proposed for it) 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 logico-mathematical or empirical concepts.
> (Carnap, 1950b, p.3, emphasis original)

Importantly, according to Carnap's conception, a successful explicatum for an inductive assumption need not be true or justified in any non-trivial sense. It must only be useful, in some cases, to replace the informal explicandum with the

formal explicatum. This can be seen from some remarks by Carnap on Russell's theory of definite descriptions, which Carnap took to be an explicative project:

> The different interpretations of descriptions are not meant as assertions about the meaning of phrases of the form the so-and-so in English, but as proposals for an interpretation and, consequently, for deductive rules, concerning descriptions in symbolic systems. Therefore, there is no theoretical issue of right or wrong between these various conceptions, but only the practical question of the convenience of the different methods.
> (Carnap, 1947a, p.33)

Due to this commitment to explication, Carnapian inductive logic not straightforwardly normative. According to Carnap, the main purpose of inductive logic was the useful formalisation of already-given inductive assumptions, rather than the demonstration that some assumptions are better than others. Carnap thought that inductive logic shared this non-normativity with deductive logic, which he took to be a formal approximation of common patterns of deductive reasoning. This view can be seen from the following quotation:

> First I wish to emphasize that inductive logic does not propose new ways of thinking, but merely to explicate old ways. It tries to make explicit certain forms of reasoning which implicitly or instinctively have always been applied both in every day life and in science. This is analogous to the situation at the beginning of deductive logic. Aristotle did not invent deductive reasoning; that had gone on as long ago as there was human language. If somebody had said to Aristotle: "What good is your new theory to us? We have done well enough without it. Why should we change our ways of thinking and accept your new invention?", he might have answered: "I do not propose new ways of thinking, I merely want to help you to do consciously and hence with greater clarity and safety from pitfalls what you have always done. I merely want to replace common sense with exact rules" It is the same with inductive logic. ...
>
> Since inductive logic merely intends to explicate common ways of inductive reasoning, the question of its usefulness leads back to the general question: Is it desirable that procedures which are generally applied, though only intuitively or instinctively, are brought into the clear daylight, analyzed and systematized in the form of exact rules? Whoever gives an affirmative answer to this general question will acknowledge the importance of the special problem of explicating inductive reasoning, that is, of constructing a system of inductive logic with rules as exact as those of the older, well-established system of deductive logic.
> (Carnap, 1953, p.189)

As we shall see in what follows, discussions of Carnapian inductive logic often

overlook this crucial aspect.

Despite the fact that Carnapian inductive logic was not normative but explicative, questions of justification can be indirectly relevant to inductive logic. One might, for example, be particularly interested in explicating inductive assumptions that one thinks are justified in a particular knowledge situation, in order to find out how someone in that situation ought to reason, given those assumptions. In such cases it is especially important to keep in mind that the normative question of whether a particular inductive assumption is justified is distinct from the inductive logical question of how best to formalise it.

2.4.2   Systems of inductive logic

A second hard core commitment of Carnapian inductive logic was to a particular kind of explicatum for inductive assumptions, namely systems of inductive logic. The following exposition of this key concept matches Carnap's on all substantive points and uses the same terms as Carnap to describe abstract objects, but follows a different presentation that is heavily informed by that of Paris and Vencovská (2015). This presentation is more succinct than Carnap's and, I hope, should also be more accessible to modern readers.

A system of inductive logic is an ordered pair $(D, \mathcal{M})$ consisting of a set D, called an 'inductive logical domain' or simply 'domain', and a set $\mathcal{M} = \{m : D \to \mathbb{R} \mid \text{axioms}\}$ of real valued 'measure functions' satisfying certain formal conditions or 'axioms'.

Inductive logical domains contain members that represent possible objects of inductive assumptions; typically these are propositions, though we shall see later that inductive assumptions can also concern states of evidence. Measure functions represent possible degrees of plausibility of such objects. Measure functions can often canonically be associated with two-place 'conditional' measure functions with the form $m_{cond} : D \times D \to \mathbb{R}$ and interpreted as representing attitudes about which propositions evidentially support which others. Carnap used the term 'confirmation function' or 'c-function' to describe conditional measure functions.

The procedure for associating unconditional and conditional place measure functions in the case for commonly used domains and axiom systems is described at (Paris and Vencovská, 2015, Ch. 4). In the typical case where D is a propositional or predicate domain, m is a probability function, $\theta, \phi \in D$ and $m(\phi) > 0$, the conditional measure function m_{cond} associated with m is defined so that $m_{cond}(\theta \mid \phi) = \frac{m(\theta \wedge \phi)}{m(\phi)}$.

I shall now describe inductive logical domains and axioms in more detail, highlighting important examples that feature in the rest of the thesis.

Inductive logical domains

Inductive logical domains are sets whose members are interpreted as representing propositions about which one might make inductive assumptions. The following two kinds of domain were studied in the most detail by Carnapian inductive logicians.

Finite Propositional domains

The simplest kind of inductive logical domain is the set of sentences SL_{prop} of a finite propositional language L_{prop}. Finite propositional languages consist of propositional variables $X_1, X_2, \ldots, X_n$ and connectives $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$, together with rules determining which combinations of these symbols are well-formed sentences.

An atomic state of a propositional language is a sentence with the form $\omega_i = \pm X_1 \wedge \ldots \wedge \pm X_n$, where each $\pm$ indicates that the following propositional variable appears either with or without a negation symbol in front of it.

Sentences of finite propositional languages are typically interpreted as representing propositions. For example, the sentence $X_1 \wedge \neg X_2$ might be interpreted as representing the proposition that xylophone 1 is in tune but xylophone 2 is not.

Unary predicate domains

In order to capture more complicated kinds of inductive reasoning, systems of inductive logic often employ a richer kind of domain: unary predicate domains. A unary predicate domain is the set SL_{pred} of a unary predicate language. Unary predicate languages consist of unary predicate symbols $P_1, P_2, \ldots, P_k$, constant symbols $a_1, a_2, \ldots, a_n$, variable symbols $x_1, x_2, \ldots$, quantifier symbols $\forall$ and $\exists$, connectives $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$ and rules specifying which combinations of these symbols are well-formed formulae and sentences.

An atom of a unary predicate language is a formula with the form $\alpha_j(x) = \pm P_1 x \pm \wedge P_2 x \pm \wedge \ldots \wedge \pm P_k x$ where, as above, each $\pm$ indicates either the presence or absence of a negation symbol. When a constant a_i is substituted for the variable x in an atom, the result is an 'atomic sentence' which, in a sense, describes a_i as fully as possible within the given language.

Terms of the form $\theta(a_1, \ldots, a_m)$ should be understood as sentences that mention all and only constants $a_1, \ldots, a_m$.

Sentences of unary predicate languages are typically interpreted so that the constant symbols represent individuals, such as experiments, and the predicate symbols represent properties of these individuals, such as possible experimental results. For example, the sentence $P_1 a_1 \wedge \neg P_5 a_2$ might be interpreted as representing the proposition that, in an experiment in which apples are sequentially

drawn out of a sack and examined, apple 1 is found to be green, while apple 2 is found not to be fresh.

Unary predicate domains can be made richer by making the set of constant symbols countably infinite, by introducing non-unary relation symbols, by introducing equality symbols or by separating predicates into 'families' and specifying that every constant must instantiate exactly one predicate in every family.

Axioms

The most important features of systems of inductive logic are the axioms which determine sets of measure functions. These are typically interpreted as representing conditions that inductive attitudes must satisfy in order to be compatible with the collection of assumptions that the overall system of inductive logic formalises.

Inductive logical axioms must be formal in the sense that they only mention measure functions' logico-mathematical properties, and not their interpretations or those of their domains. Carnap discusses this point at (Carnap, 1950b, § 43). For example, it is possible for an inductive logical axiom to excludes measure functions which assign positive numbers to logical contradictions, since whether or not a given measure function satisfies this condition depends only on its logico-mathematical properties. On the other hand, an inductive logical axiom could not exclude measure functions that assign positive numbers to false sentences. Whether or not a measure function assigns positive values to false sentences depends on which sentences are false: this is not a logico-mathematical property of the measure function but rather a feature of the interpretation of its domain.

Some important inductive logical axioms are as follows. All of them assume a unary predicate domain SL_{pred}. Many interesting facts about these axioms and more are demonstrated in Paris and Vencovská (2015).

Probability (unary predicate domains)

The probability axioms formalise the inductive assumption that attitudes about propositions' plausibilities can be represented by probability functions. It is important to note that representing this inductive assumption using the probability axioms need not imply any commitment to the epistemological view that states of rational partial belief should be representable by probability functions, which we will discuss in chapter 4.

> Axioms of probability (unary predicate domains): **Prob$_{\mathbf{pred}}$**
> For any sentences θ and $\phi \in SL_{pred}$, $m : SL_{pred} \to [0,1]$ must be such that:
>
> > $P1_{pred}$ If θ is logically true then $m(\theta) = 1$.

$P2_{pred}$ If $\neg(\theta \wedge \phi)$ is a logical truth, then $m(\theta \vee \phi) = m(\theta) + m(\phi)$.

$P3_{pred}$ For any sentence with the form $\exists x\theta(x)$, $m(\exists x\theta(x)) = \lim_{n \to \infty} m(\theta(a_1) \vee \theta(a_2) \vee \ldots \vee \theta(a_n))$

These conditions ensure that measure functions are probability functions. The first two conditions ensure that, as probability functions, measure functions satisfying them give maximum values to logical truths and that the probabilities of mutually exclusive sentences are additive. The third condition $P3_{pred}$ allows the probability of sentences with quantifiers to be determined by those of quantifier-free sentences. This axiom preserves the intuitive meaning of quantifiers, so long as the individuals represented by constant symbols are assumed to exhaust the relevant universe of discourse. See (Paris and Vencovská, 2015, Ch. 3) for more on this point.

Constant exchangeability

This axiom attempts to formalise the inductive assumption that the identity of particular individuals makes no difference to the plausibility of propositions about them.

Axiom of constant exchangeability: **Ex**
For any tuples of constants $(a_1, \ldots, a_l)$ and $(a_{\sigma(1)}, \ldots, a_{\sigma(l)})$, where $\sigma : \mathbb{N}^+ \to \mathbb{N}^+$ is a permutation, and any sentence $\theta(a_1, \ldots, a_l) \in SL_{pred}$, $m(\theta(a_1, \ldots, a_l)) = m(\theta(a_{\sigma(1)}, \ldots, a_{\sigma(l)}))$.

See (Paris and Vencovská, 2015, Ch. 6-9) for a discussion of the mathematical properties of constant exchangeability and its relationships with other important axioms.

Instantial relevance

Axioms of instantial relevance formalise inductive assumptions according to which, when experimental results of a certain kind occur, the plausibility of other results of the same kind should be affected in a certain way. The most common such axioms are as follows.

Axiom of non-negative instantial relevance: **PIR**
For any atom $\alpha_j(x)$ of L_{pred} and sentence $\theta(a_1, \ldots, a_l) \in SL_{pred}$,
$m(\alpha_j(a_{l+2}) \mid \alpha_j(a_{l+1}) \wedge \theta(a_1, \ldots, a_l)) \geq m(\alpha_j(a_{l+2}) \mid \theta(a_1, \ldots, a_l))$

Axiom of postive instantial relevance: **IR$^+$**
For any atom $\alpha_j(x)$ of L_{pred} and sentence $\theta(a_1, \ldots, a_l) \in SL_{pred}$,
$m(\alpha_j(a_{l+2}) \mid \alpha_j(a_{l+1}) \wedge \theta(a_1, \ldots, a_l)) > m(\alpha_j(a_{l+2}) \mid \theta(a_1, \ldots, a_l))$

Axiom of negative instantial relevance: **IR$^-$**
For any atom $\alpha_j(x)$ of L_{pred} and sentence $\theta(a_1, \ldots, a_l) \in SL_{pred}$,
$m(\alpha_j(a_{l+2}) \mid \alpha_j(a_{l+1}) \wedge \theta(a_1, \ldots, a_l)) < m(\alpha_j(a_{l+2}) \mid \theta(a_1, \ldots, a_l))$

In each case, constants are interpreted as representing experiments and the atom α_j as representing a certain kind of result. **PIR** allows only measure functions representing attitudes according to which the proposition that a result of a certain kind occurred does not count as evidence against the proposition that another result of the same kind occurs subsequently. **IR**$^+$ represents a stronger assumption according to which instances must positively count in favour of subsequent instances of the same kind. **IR**$^-$ represents an assumption that goes the other way and insists that experimental results of a certain kind make subsequent results of the same kind strictly *less* plausible.

The axiom of non-negative instantial relevance is sometimes called the 'principle of instantial relevance', hence the abbreviation **PIR** and has been the subject of the most study. See, for example, (Paris and Vencovská, 2015, Ch. 11). It is helpful to remember, however, that other axioms of instantial relevance exist.

Johnson's sufficientness postulate
This axiom attempts to formalise the inductive assumption that, in order to determine how plausible a sample of n experimental results makes an as-yet unobserved experimental result of a certain kind, the only relevant information is the number n of experiments that constitute the sample, together with the number r of results of the same kind as the one whose future plausibility is being assessed.

> Johnson's sufficientness postulate: **JSP**
> $m(\alpha_j(a_{n+1}) \mid \bigwedge_{i=1}^{n} \alpha_{h_i}(a_i))$ depends only on n and $r = \mid \{i : h_i = j\} \mid$

In this formulation the atoms of L_{pred} are assumed to represent possible results of experiments and the constant symbols to represent experiments. The number r is the number of constant symbols in $\bigwedge_{i=1}^{n} \alpha_{h_i}(a_i)$ that instantiate the atom α_j, representing the number of experiments that had the result represented by α_j.

Johnson's sufficientness postulate was first studied by W.E. Johnson, who characterised it as a 'postulate' rather than an axiom, explaining the first and third part of its name, which, according to (Paris and Vencovská, 2015, p. 103), was coined by I.J. Good. It is called a 'sufficientness' postulate because, according to the assumption that it formalises, the numbers n and r are sufficient statistics.

Example of a system of inductive logic

An example of a system of inductive logic is $(SL_{prop}, \mathcal{M}_{\textbf{prob}})$. This system of inductive logic consists of the set of sentences SL_{prop} of a propositional language L_{prop}, together with the set of all probability functions on those sentences. This set of measure functions is determined by the probability axioms for propositional domains, which are as follows:

Probability (propositional domains)

> Axioms of probability (propositional domains)
> For any sentences θ and $\phi \in SL_{prop}$, $m : SL_{prop} \to [0,1]$ must be such that:
>
> > $P1_{prop}$ If θ is logically true then $m(\theta) = 1$.
> >
> > $P2_{prop}$ If $\neg(\theta \wedge \phi)$ is a logical truth, then $m(\theta \vee \phi) = m(\theta) + m(\phi)$.

2.4.3 Tolerance

The final key commitment of Carnapian inductive logic was to the principle that the success or failure of inductive logical explicatations should not be judged on philosophical grounds, but rather on pragmatic ones. There were two aspects to this tolerant stance. On the one hand, Carnap had a general view, that went beyond the particular case of inductive logic, according to which 'external questions' about whether to accept or reject the use of logical systems for particular purposes should be answered by pragmatic rather than philosophical means. Secondly, within inductive logic Carnap was careful to distinguish what he called questions of 'pure inductive logic', in which only the formal properties of systems of inductive logic are taken into consideration, from questions of 'applied inductive logic', for the purposes of which the components may be thought of as representing features of the non-formal world, like propositions and inductive assumptions. Carnap thought that all questions as to how well systems of inductive logic perform as explicata should be considered within the second category.

Pure and applied inductive logic

Carnapian inductive logic had two distinct branches: pure and applied inductive logic. This distinction is explained most clearly at (Carnap, 1971a, § 4) but early Carnap made it as well, as can be seen from (Carnap, 1950b, §44 A), where Carnap distinguishes 'logical' and 'methodological' questions in very similar terms to those of (Carnap, 1971a, § 4).

Pure inductive logic addressed questions about systems of inductive logic that could be answered without interpreting their domains. It therefore investigated only their mathematical properties, and did not address the question of justification. Applied inductive logic dealt with all tasks that required systems of inductive logic to be interpreted. In particular, the task of discovering which systems of inductive logic can usefully formalise which inductive assumptions required that the systems be interpreted, and so fell within the purview of applied inductive logic.

When he discusses the proper roles of pure and applied inductive logic at (Carnap, 1971a, § 4), Carnap draws a parallel with pure and applied geometry.

> The relation between pure and applied IL is somewhat similar to that between pure (mathematical) and empirical (physical) geometry. . . . In mathematical geometry we speak abstractly about certain numerical magnitudes of geometrical entities, for example, the lengths of the sides of a triangle and the measures of the three angles, but without specifying a procedure of measuring these magnitudes. General theorems are given, stating mathematical relations between these parameters. A variety of possible structures of three-dimensional space, the Euclidean and various non-Euclidean structures, are systematically studied. But the question which of these possible structures is the actual structure of the space of nature is not even raised. This question belongs to physical geometry.
> (Carnap, 1971a, p. 69)

Carnap refers his reader to Hempel (1945) for a fuller discussion of the distinction between pure and applied geometry that he has in mind. It is worth considering Hempel's remarks on pure geometry:

> Geometry thus construed is a purely formal discipline, we shall refer to it also as *pure geometry*. A pure geometry, then—no matter whether it is of the euclidean or of a non-euclidean variety—deals with no specific subject-matter; in particular, it asserts nothing about physical space. All its theorems are analytic and thus true with certainty precisely because they are devoid of factual content. Thus, to characterize the import of pure geometry, we might use the standard form of a movie-disclaimer: No portrayal of the characteristics of geometrical figures or of the spatial properties or relationships of actual physical bodies is intended, and any similarities between the primitive concepts and their customary geometrical connotations are purely coincidental.
> (Hempel, 1945, p.12, emphasis original)

With the help of this quotation from Hempel we can see that the analogy that Carnap wished to draw between pure geometry and pure inductive logic is as follows. Carnap thought that, just as the proper role of pure geometry was not to indicate which structure should be used to model physical space, so too it was beyond the role of pure inductive logic to specify which system of inductive logic should be used for a given purpose.

The following quotation confirms Carnap's view that the evaluation of systems of inductive logic should only proceed within applied inductive logic:

> Justifying an inductive method and, more specifically, offering reasons for the acceptance of a proposed axiom, is a kind of reasoning that lies outside pure IL and takes into consideration the *application* of c-functions. (Carnap, 1971a, p. 105, emphasis original)

This aspect of Carnap's use of the term 'pure inductive logic' contrasts with that of the contemporary Manchester school of inductive logicians, whose approach is more liberal. The Manchester school devote particular attention to the question of which inductive logical axioms formalise rationality constraints that bind agents who find themselves in what I call the 'Manchester blank slate situation'. We shall discuss blank slate situations in detail in section 2.9.1 below. Briefly, the Manchester blank slate situation occurs when an agent finds themself required to attach degrees of plausibility to members of an inductive logical domain without knowing how these members are associated with propositions. According to the Manchester school, axioms that formalise blank-slate rationality constraints constitute the subject matter of 'pure inductive logic'. Consequently, as the Manchester school use the term, 'pure inductive logic' includes activities—identifying blank-slate rationality constraints—that would be ruled out according to Carnap's use of the term.

Carnap's general views about logical systems

Carnap's views about 'external questions' whether to accept or reject a logical system are expressed in Carnap (1950a) as follows:

> . . . we take the position that the introduction of the new ways of speaking does not need any theoretical justification because it does not imply any assertion of reality. . . .
>
> Above all, [the introduction of a new system] must not be interpreted as referring to an assumption, belief, or assertion of "the reality of the entities." There is no such assertion. An alleged statement of the reality of the system of entities is a pseudo-statement without cognitive content. To be sure, we have to face at this point an important question; but it is a practical, not a theoretical question; it is the question of whether or not to accept the new linguistic forms. The acceptance cannot be judged as being either true or false because it is not an assertion. It can only be judged as being more or less expedient, fruitful, conducive to the aim for which the language is intended.
> (Carnap, 1950a, § 3)

Semantic rules, Carnap thought, cannot be true or false, but only more or less useful. Consequently the axioms that determine systems of inductive logic, which Carnap saw as semantic rules, cannot be justified on a theoretical basis, but only on the basis of how useful they are for particular purposes. This connection is made by Carnap himself in the following quotation:

> I agree with Burks' view that. . . questions of justification of induction in general. . . are external questions.
> (Carnap, 1963a, p. 981-2)

2.4.4   Summary of Carnapian inductive logic's hard core

To sum up, the following commitments can be thought of as comprising the hard core of Carnapian inductive logic:

$HC1$ **Explication** Inductive assumptions can usefully be explicated.

$HC2$ **Systems of inductive logic** Systems of inductive logic can be good explicata for inductive assumptions.

$HC3$ **Tolerance** All systems of inductive logic should be tolerated at a philosophical level: systems of inductive logic should only be evaluated on pragmatic grounds, in the context of a given interpretation and application.

As mentioned above, Lakatos tended to conceive of hard-core commitments as truth-directed claims about the nature of the world, whereas $HC1-3$ are more akin to methodological prescriptions. Nonetheless, $HC1-3$ play exactly the role required by Lakatos's methodology. They were all advocated by Carnapian inductive logicians at all stages of their research programme's development. They are constitutive of the research programme: if a research programme does not explicate inductive assumptions, uses means other than systems of inductive logic to explicate them or seeks to justify systems of inductive logic on theoretical rather than practical grounds, then it is certainly not Carnapian inductive logic. *HC1-3* are also specific enough to distinguish Carnapian inductive logic from its rivals. For example, Timothy Williamson's conception of evidential probability, as presented at (Williamson, 2002, p.209-237), denies *HC3* by postulating true degrees of evidential support and claiming that they satisfy the probability axioms. Karl Popper, for example in (Popper, 1959, Ch. 1), denied $HC1$, arguing that inductive assumptions are irrelevant to science and therefore unworthy of explication. Putnam (1963) rejected $HC2$, arguing instead that inductive assumptions should be formalised using different mathematical objects.

2.5   Other components

According to the methodology of scientific research programmes, the main components of a research programme, besides the hard-core of unalterable commitments, are the protective belt of secondary commitments, the negative heuristic and the positive heuristic. In the case of Carnapian inductive logic, I claim that these were as follows.

2.5.1 Protective belt

Carnapian inductive logic's protective belt consists of auxiliary hypotheses, each of which asserts that a particular system of inductive logic successfully explicates a particular inductive assumption. Since, in order to be successful, an explicatum must both approximate the thing that it formalises and be useful, claims in Carnapian inductive logic's protective belt have at least the following components. On the one hand they assert that a certain system of inductive logic adequately approximates a certain inductive assumption; on the other they claim that that system is useful for certain purposes.

For example, in Carnap (1952a), Carnap argued for the explicative merits of a certain system of inductive logic, namely $(SL_{pred}, \mathcal{M}_{1952})$ where $\mathcal{M}_{1952}$ contains all and only the measure functions in the so-called 'continuum of inductive methods'. Carnap claimed that this system successfully explicates a commonly-made collection of inductive assumptions. First he argued, for each of eleven axioms with which he characterised $\mathcal{M}_{1952}$[5], that it approximates a feature of common-sense inductive reasoning. He then attempted to show that $(SL_{pred}, \mathcal{M}_{1952})$ is useful by finding applications for it to problems in statistics, such as the debate over whether estimators should be unbiased. In order to make $(SL_{pred}, \mathcal{M}_{1952})$ easier to apply in this way, Carnap showed that the measure functions in $\mathcal{M}_{1952}$ can be described in terms of a single real-valued parameter λ.

It is important to recognise that, although claims as to the success of failure of particular explicata are prominent in Carnap's writing on inductive logic, Carnapian inductive logic did not make a hard core committment to any such claim. As we have seen, Carnap was in fact committed to a tolerant approach to all such applied inductive logical questions. It is therefore, I believe, unwise to identify Carnpian inductive logic with any particular system of inductive logic or axiom.

2.5.2 Negative heuristic

A research programme's negative heuristic instructs researchers as to how they ought to react to 'anomalies': inconvenient facts that sit awkwardly with its hard core commitments. In the case of Carnapian inductive logic, anomalies are cases of inductive reasoning that prove difficult to explicate using currently available systems of inductive logic.

For example, (Goodman, 1946, p.383) asks us to consider a situation where a sequence of observations displays a periodic pattern. Perhaps an observation of a red object is invariably followed immediately by an observation of a black object. A natural inductive assumption asserts that, after many repetitions

---

[5]Carnap and his co-workers later improved upon this characterisation, successfully defining $\mathcal{M}_{1952}$ using only the axioms **Prob**, **Ex** and **JSP**. See (Paris and Vencovská, 2015, Ch. 16) for more on Carnap's continuum.

of this pattern, observing a red object ought to make observing a black object at the next stage highly plausible. However, this assumption is difficult to capture without abandoning the technically important axiom of constant exchangeability which, according to a standard interpretation, forces the order of observations to be evidentially irrelevant. Systems of inductive logic that do not include constant exchangeability are not as well-understood as those that do include it: Goodman's example therefore seems to constitute an anomaly.

Carnapian inductive logic had two ways to accommodate such anomalies without surrendering any hard core dogmas. The first blames the crude state of its current machinery. Just as Newton's early models were not sophisticated enough to account for many celestial observations, so, Carnapian inductive logic's negative heuristic might say, it would be unreasonable to expect early inductive logic to account for all inductive assumptions. Carnap used this strategy in Carnap (1948), his original response to Goodman, arguing that more sophisticated systems of inductive logic would be needed to account for reasoning based on periodicities.

Alternatively, failures of explication can be explained by directing attention towards the explicandum. Perhaps a particular system of inductive logic does not explicate some inductive reasoning very well because certain features of the reasoning—perhaps distinctive background information or tacit assumptions— have not yet been accounted for. For example, applying inductive logic to the anomalous situation described above makes clear that the order of observations is evidentially relevant in that situation. By explaining anomalies in this second way, what seems initially to be a failure of inductive logic can turn out to be a success, as difficulties with explication lead to new discoveries about our inductive practices. Lakatos was particularly fond of this kind of application of the negative heuristic: at (Lakatos, 1968b, p.169-170) he extolls the ability of Newtonian researchers to "[turn] each new difficulty into a new victory of their programme.".

2.5.3   Positive heuristic

A research programme's positive heuristic instructs researchers as to how they should expand their research programme's protective belt, making steadily more secondary commitments. Broadly speaking, Carnapian inductive logic's positive heuristic instructed researchers to do so by making more and more claims according to which systems of inductive logic successfully explicate inductive assumptions.

Carnap expressed more specific plans for expansion in Carnap (1945a) and (Carnap, 1950b, § 49). In Carnap (1945a) he set out different kinds of inductive assumptions that he would like to be able to capture using inductive logic, including assumptions that take into account the similarity or dissimilarity of individuals and assumptions based on general laws. At (Carnap, 1950b, § 49)

he envisaged the construction of more sophisticated domains, such as the sets of sentences of a polyadic language. Carnap aimed eventually to develop systems of inductive logic that would be rich enough to allow a full array of propositions and inductive assumptions to do with real-valued magnitudes to be represented.

Although Carnap did not lay out an explicit programme for expansion in this respect, the development of new axioms representing previously un-explicated inductive attitudes and the investigation of such axioms' mathematical properties was clearly another goal. The presence of this goal can be inferred from the systematic way that Carnap and his coworkers carried it out.

2.6 How Carnapian inductive logic developed

Lakatos's evaluative criteria are diachronic: rather than classifying a research programme as progressive or degenerate based on its properties at a particular moment in time, its assessment depends on how it develops as time goes by. In order to evaluate Carnapian inductive logic according to Lakatos's criteria, then, we will need to assess how it changed during the course of its existence.

In this section I document the main ways in which Carnapian inductive logic changed: this will make it possible to carry out a diachronic assessment.

2.6.1 Less restrictive systems of inductive logic

Early Carnapian inductive logic, up to and including Carnap (1950b), focuses exclusively on systems of inductive logic where the set of measure functions contains only one measure function. As Carnapian inductive logic developed, it started to investigate systems of inductive logic with steadily larger sets of measure functions. The set $\mathcal{M}_{1950}$ of measure functions allowed by the system that Carnap investigated in Carnap (1950b) contains only the single measure function $m^\star$. In Carnap (1952a) the focus is on a set $\mathcal{M}_{1952}$ that contains uncountably infinitely many measure functions, including $m^\star$. Carnap (1971a) and Carnap (1980) focus on systems with a still larger set $\mathcal{M}_{1971}$ of which $\mathcal{M}_{1952}$ is a proper subset.

The reason for this steady liberalisation seems to have been technical progress. Carnap restricted his attention to sets of measure functions whose members can be fully described by relatively few real-valued parameters, which have natural interpretations. For example, each measure function in the set $\mathcal{M}_{1952}$ is fully described by the real-valued parameter $\lambda \in [0, \infty]$, which can naturally be interpreted as representing sensitivity to observational evidence.

As Carnapian inductive logic developed, Carnap and his co-workers discovered more sophisticated parameterisations that could describe more comprehensive sets of measure functions. Consequently, their attention shifted towards the

task of investigating and finding uses for the larger sets of measure functions they could describe using their new tools.

2.6.2 Different interpretative requirements

Early applied Carnapian inductive logic featured several requirements stipulating how logical languages featuring in systems of inductive logic ought to be interpreted. These requirements included, for example, the stipulation that logical constants be interpreted as representing distinct individuals and that predicates represent possible properties of those individuals. Another requirement, the 'requirement of independence' (Carnap, 1950b, §18 B), required that inductive logical languages should be interpreted so that all distinct sentences with the form $P_i(a_k)$ are logically independent.

Late Carnapian inductive logic relabelled all requirements of this kind 'basic assumptions' and developed a framework for expressing them explicitly as metalinguistic stipulations, so that different requirements could easily be chosen for the purposes of different applications. For example, in an application of inductive logic to the investigation of an urn full of coloured balls, a basic assumption might express the fact that each ball must have one and only one colour by stipulating that each constant of a language must instantiate exactly one predicate. See (Carnap, 1971a, Ch. 5) for Carnap's discussion of basic assumptions.

2.6.3 Families of predicates

Later Carnapian inductive logic discussed logical predicates slightly differently from early Carnapian inductive logic, organising them into 'families', corresponding to mutually exclusive and exhaustive categories of properties such as possible colours, weights and volumes that a sampled individual might have. It was stipulated by means of a basic assumption that every constant instantiates exactly one member of each family. Families are introduced at (Carnap, 1971a, p. 43-47), and the salient basic assumption is set out at (Carnap, 1971a, p. 81).

This additional structure had several attractive features. It allowed impossibilities like the whole surface of a ball being both black and red to be ruled out and brought inductive logic closer to mainstream statistics, since predicates with a family structure can be thought of as values of random variables.

While families of predicates, and the means of describing them logically, feature prominently in late Carnapian inductive logic, their introduction did not fundamentally alter the research programme. The organisation of predicates into families is simply a matter of labelling, which could just as well have been executed in early Carnapian inductive logic. The basic assumption stipulating that individuals instantiate one predicate in each family, on the other hand, was novel. However, it is properly seen as a generalisation of the original approach

rather than a departure, because the presentation without a family structure can be recovered by stipulating that all predicates belong to different families.

2.6.4 Measure-theoretic reformulation

Carnap recognised that the formal languages that featured in his early work on inductive logic were unfamiliar to most practising scientists. In order to remedy this situation, later Carnapian inductive logic discussed systems of inductive logic whose domains are algebras generated by an underlying state space rather than sets of sentences of a logical language.

This change was cosmetic, as the axioms that Carnap and his co-workers investigated did not exploit the differences between these kinds of domains. Instead they simply reformulated axioms on sentence-domained measure functions in the new measure-theoretic framework.

2.6.5 Versatility

A final noteworthy change concerns the versatility of systems of inductive logic. Early Carnap seems to have believed that systems of inductive logic ought to be highly versatile, so that, having chosen a particular system for a particular task, an inductive logician should generally use the same system for all other applications of inductive logic. This is clear from Carnap's postulation in Carnap (1950b) of the generally accepted set of inductive assumptions 'probability$_1$', which is discussed in detail below in section 2.11.3, and from the following passage. Another indication of this position comes from the following passage, where early Carnap argues that it is better to choose systems of inductive logic that are generally applicable, rather than choosing different systems for the purposes of different applications, or "inductive problems":

> Suppose that X has chosen a certain inductive method ['inductive method' is Carnap's word for a system of inductive logic with one measure function $m \in \mathcal{M}_{1952}$ from his continuum.] and used it during a certain period for the inductive problems which occurred. If, in view of the service it has given him, he is not satisfied with it, he may at any time abandon it and adopt a method that seems to him preferable. This is not the same as a change in method from problem to problem. Once he adopts the new method, he will apply it to all inductive problems... One inductive method is here envisaged as covering all inductive problems.
> (Carnap, 1952a, p.54, square parentheses added)

Later Carnap took a more liberal view on this applied inductive logical question, allowing different systems of inductive logic to be chosen to suit particular

applications, or in Carnap's words 'investigations'. The first indication of this change that I am aware of comes from correspondence with Shimony:

> I emphasize now that it seems best to regard inductive logic as always applied to a particular scientific investigation, and therefore as referring to the conceptual system (the "language") of that investigation.
>
> (Shimony and Carnap, 1969, parentheses original)

In this correspondence Carnap agrees with Shimony that that different applications of inductive logic may require both different domains and different sets of measure functions and assures him that this change would be put into practice in his forthcoming book on the topic.

This did in fact happen: at (Carnap, 1971a, p.44) Carnap explicitly allowed that features of inductive logical domain relevant to a particular application may be chosen on the basis of the appropriate amount of precision in a particular case. In an example Carnap argued that it might be worthwhile to use predicates that represent individuals' ages in whole years rather than, say, in months, "because we decide that for the purposes of our statistical investigation higher precision is unnecessary".

Later in the same work Carnap allows that different applications may make different axioms appropriate. For example, depending upon the application, the 'principle of symmetry', which we have called 'constant exchangeability', may or may not be appropriate:

> In some series of events, the temporal order has no influence. For example, we find that, in a long series of throws of any die, even if it is loaded, the relative frequency of a an ace among those throws that follow immediately on an ace (or a deuce), is not essentially different from the relative frequency of an ace in the whole series. Thus experience shows here statistical independence; and therefore we treat the members of such a series as inductively independent. In contrast, we find a strong dependence in a series of meteorological observations made at a fixed place at noon every day... The relative frequency of rain immediately following on rain is much higher than in the whole. If now in elementary inductive logic, we decide to accept the principle of symmetry *for a certain investigation*, this decision may either be based on the assumption that statistical independence actually holds, or else on the assumption that there is some dependence, but so small that we may neglect it for the sake of simplicity...
>
> (Carnap, 1971a, p. 120 emphasis added)

In sum, it seems that, as time went by, Carnap changed his view on the applied inductive logical question of how versatile systems of inductive logic should be. At first he thought that systems of inductive logic should be very versatile, whereas later he thought that they should be less versatile.

2.7 Carnapian inductive logic was progressive

I shall now explain what it would have meant for Carnapian inductive logic to exhibit the characteristics that, according to Lakatos's methodology, made a research programme progressive or degenerate. I consider each of the characteristics that Lakatos saw as signs of progress—theoretical progress, empirical progress and heuristic power—and claim that, to varying extents, Carnapian inductive logic exhibited all three.

Theoretical progress

Carnapian inductive logic would display theoretical progress if its researchers managed to add to its protective belt, making more and more claims to have successfully explicated inductive assumptions.

The researchers who carried out Carnapian inductive logic seem to have done this. Firstly, they fulfilled their positive heuristic's demand to investigate more complex domains. Carnap documents some tentative steps in this direction at Carnap (1951) and later at (Carnap, 1971a, § 2B). While Carnap and his co-workers did not make very substantial progress in this regard, subsequent researchers have done much better. The Manchester school of inductive logicians has developed and extended the investigation of polyadic languages—see for example Landes et al. (2008) and Landes (2009). Blackwell and MacQueen (1973) extend a broadly Carnapian approach to domains that can represent continuous magnitudes.

The second proposed direction of expansion was more thoroughly fulfilled in Carnap's lifetime. Carnap and his co-workers devised various novel axioms to capture previously un-described inductive phenomena, such as long-run agreement with relative frequency (Carnap, 1980, § 20A), symmetry among properties (Carnap, 1980, p.75-76), degree of resistance to changing one's mind (Carnap (1952a)), inductively relevant similarity between properties and individuals (Carnap, 1980, Ch. 16) and uncertain evidence (Carnap (1967b)). This project has also been continued by subsequent inductive logicians: for example, Hill (2013) investigates new principles of analogical reasoning.

Carnapian inductive logic's protective belt also expanded in ways that were not anticipated by its positive heuristic. The change from rigid interpretative requirements to flexible basic assumptions allowed the explication of inductive assumptions involving logical dependencies, and therefore represented an expansion of the protective belt. However, at (Carnap, 1950b, § 18), early Carnap describes the rigid requirements as "essential for inductive logic" - clearly this expansion was not anticipated. Similarly, as we have seen, the shifts to less versatile systems of inductive logic and larger sets of measure functions allowed more inductive assumptions to be explicated, but represented minor departures from Carnap's earlier approach.

Carnapian inductive logic therefore seems to have made considerable theoretical progress.

Empirical progress

Since Carnapian inductive logic did not aim to produce novel predictions about the world, it is not obvious what would constitute empirical progress. Nonetheless, it is possible to find an appropriate criterion by considering the role that empirical progress plays in the history of science, on Lakatos's account - that of distinguishing successful research programmes and providing an ultimate goal for researchers. Whereas, according to Lakatos, the ultimate goal of any scientific research programme is to produce novel, verified, predictions, the ultimate goal of Carnapian inductive logic was to produce useful formalisations. A plausible definition of empirical progress for the case of Carnapian inductive logic would therefore require the explicata that it produces to prove useful in science and everyday life.

It must be admitted that there are few examples of scientists fruitfully explicating inductive reasoning using Carnapian inductive logic. Carnap (1952a) made several points that were relevant to and consciously directed towards contemporary debates among practising statisticians, but, as Carnap lamented in his intellectual autobiography, these gained little attention from their target audience:

> I had expected that statisticians would offer counter-arguments to defend their customary methods; but so far I have seen none. I have often noticed in discussions that it seems very difficult for those who have worked for years within the given framework of mathematical statistics to adapt their thinking to the unfamiliar concept of logical probability.
> (Carnap, 1963d, p. 76)

Skyrms (1996) argues that a strand of modern statistics is similar in spirit to Carnapian inductive logic, but does not demonstrate any direct influence. At least as far as proving useful in science goes, then, Carnapian inductive logic has not yet displayed significant empirical progress.

However, there is good reason to expect that Carnapian inductive logic will be empirically progressive in the future. One reason for optimism in this regard comes from Zabell (2011). Zabell argues that Carnapian inductive logic can be used to help statisticians to choose between classes of statistical models. Many popular families of statistical models can be interpreted as systems of inductive logic: in such cases it makes sense to evaluate the models based on whether their corresponding axioms explicate appropriate inductive assumptions. For example, statistical models featuring symmetric Dirichlet priors over independent, identically distributed random variables can be evaluated according to whether Johnson's sufficientness postulate, the corresponding axiom, represents

an appropriate inductive stance. I shall develop Zabell's argument in detail in chapter 6.

Carnapian inductive logic has many other potential applications - any endeavour that involves inductive reasoning can potentially benefit from Carnap's clear philosophical account of how it can be formalised. In chapter 7 I highlight potential applications to debates over imprecise models of rational partial belief, climate models and objective Bayesian epistemology.

Lakatos allowed that successful research programmes may take a long time before exhibiting empirical progress, and need only do so intermittently. Even Newtonian gravity, he notes at (Lakatos, 1968b, p.172), needed to develop for more than a decade before Newton was prepared to "look more anxiously at the facts". Since it has the potential to be useful in science, Carnapian inductive logic cannot be dismissed as non-progressive on the basis of a lack of empirical progress.

Heuristic power

Carnapian inductive logic would show heuristic power if its protective belt tended to change in a way that was not 'ad hoc'. Lakatos explains at (Lakatos, 1970, p.125) that a change to a research programme's protective belt can be ad hoc in three ways:

- Failing to add genuine empirical content ('ad hoc$_1$').

- Merely reproducing the work of other research programmes ('ad hoc$_2$').

- Straying from the spirit of the programme's positive and negative heuristics ('ad hoc$_3$').

Adding genuine empirical content

Additions to Carnapian inductive logic's protective belt were claims according to which systems of inductive logic explicate inductive assumptions. The epistemological status of such claims is somewhat ambiguous because successful explication requires usefulness: testing whether an inductive assumption has been explicated requires testing whether its formal replacement is useful. This kind of test is less epistemologically straightforward than, for example, testing whether a planet is orbiting the sun. It is therefore difficult to say precisely how much empirical content is contained in a given addition to Carnapian inductive logic's protective belt.

However, in order to be ad-hoc in Lakatos's first sense, Carnapian inductive logic's claims to explicate inductive assumptions would need to be completely devoid of empirical content. This is surely not the case: any claim that a particular system of inductive logic explicates a particular inductive assumption has at least some empirically testable consequences. Thus it is difficult to argue

that Carnapian inductive logic made changes that were ad-hoc in Lakatos's first sense.

Reproducing other research programmes' work

Some additions to Carnapian inductive logic's protective belt were ad-hoc in Lakatos's second sense, as they reproduced the work of rival research programmes. For example, Carnapian inductive logicians explicitly made use of De Finetti's results concerning the axiom constant exchangeability—see, for example (Jeffrey, 1971, § 10) and Gaifman (1971)—as well as, seemingly unknowingly, reproducing some work by W.E. Johnson: see Zabell (1982) for details.

However, these ad hoc additions are more than counterbalanced by many additions that were not anticipated by rival research programmes. Examples of novel additions include Carnap's claim in Carnap (1952a) to have explicated degree of unwillingness to change one's mind using the parameter λ and his claim to have explicated reasoning based on similarity between properties in (Carnap, 1980, Ch. 16). Overall, the way that Carnapian inductive logic's protective belt expanded seems not to have been ad hoc in Lakatos's second sense.

The spirit of the heuristics

Finally, the development of Carnapian inductive logic stayed largely within the spirit of both its positive and negative heuristic. Although some additions to the protective belt required unforeseen changes, these changes were comparatively minor, and many other changes were anticipated.

With respect to the negative heuristic, Carnap and his coworkers managed to defend their work against perceived anomalies without sacrificing any hard-core commitments. This consistency can be seen from the preface to the second edition of *Logical Foundations of Probability*, where Carnap writes that "the major features of my theory, as explained in this book, are still maintained today" (Carnap, 1950b, p. xiii). None of the changes outlined above in section 2.6 involved rejecting any of the hard core dogmas. Carnapian inductive logicians therefore obeyed its negative heuristic.

Carnapian inductive logic also developed within the spirit of its positive heuristic, explicating steadily more inductive phenomena as it developed. In some cases later work explicitly developed earlier work, showing continuity of purpose and spirit that would surely have impressed Lakatos. For example, in (Carnap, 1945a, § 11), Carnap discusses reasoning based on similarity, arguing that this could be formalised within the system of inductive logic that he proposed in that paper. In (Carnap, 1980, Ch. 16) this kind of reasoning is discussed in considerably greater depth, and formalised using a more sophisticated system of inductive logic. While several initial goals, such as investigating polyadic languages and formalising reasoning involving continuous magnitudes, were not fulfilled within Carnap's lifetime, the development of Carnapian inductive logic

proceeded roughly as he and his co-workers might have predicted at the outset. In short, the positive heuristic was largely obeyed.

It seems that Carnapian inductive logic stayed broadly within the spirit of both its positive and negative heuristic: it was therefore not ad-hoc in Lakatos's third sense.

Summary

In sum, then, Carnapian inductive logic displayed substantial theoretical progress, adequate heuristic power and shows signs of producing empirical progress in the future. This, I claim, is sufficient for it to be considered a progressive research programme, according to Lakatos's methodology.

2.8 Discussion

Let us now consider the argument up to this point. I have introduced Carnapian inductive logic by construing it as a research programme. I classified its components according to Lakatos's methodology and set out the main ways in which it changed as it developed. I then argued that Carnapian inductive logic was progressive according to Lakatos's methodology of scientific research programmes.

If this chapter has done its job, then its novel textual arguments will have stimulated the reader's appetite for a historical reappraisal of Carnapian inductive logic, while its application of Lakatos's methodology might perhaps have made them curious about the possibility of substantive reappraisal. Committed adherents of Lakatos's evaluative framework may even be ready for both courses. However, others may retain some scepticism about Carnapian inductive logic's continuing relevance, as the methodology of scientific research programmes has several important limitations.

Subjectivity

For one thing, Lakatos's criteria are not universally accepted. One issue, raised forcefully in (Feyerabend, 1981, Ch. 10), is that many of the various steps required in order to apply the methodology of scientific research programmes are somewhat subjective. For example, practising scientists rarely explicitly distinguish between inalienable hard-core commitments and defeasible secondary ones: the existence of such a distinction must usually be inferred, using the interpreter's judgement. To the extent that such judgements are unavoidably subjective, the resultant attributions of progressiveness are correspondingly arbitrary. Feyerabend believed that, due to this issue, Lakatos fell short of providing an objective means of assessing research programmes.

Those who agree with Feyerabend's assessment of the methodology of scientific research programmes will neither be surprised nor impressed by the fact that the methodology can be applied in such a way as to make Carnapian inductive logic progressive.

Rigidity

(Laudan, 1978, p.76-78) criticises the methodology of scientific research programmes for a different reason. According to Laudan, it asks too much of scientific research programmes by insisting that they be rigid in a certain sense.

Once a research programme's positive and negative heuristic are set, Laudan notes, they are, according to Lakatos, fixed forever. No new hard-core commitments can be introduced, nor new strategies for dealing with uncooperative facts, expanding the protective belt or making novel predictions. Thus Lakatos's methodology seems to favour rigid, dogmatic research programmes over flexible ones. One might question whether this kind of rigidity is really a genuine mark of scientific progress.

The reader might suspect that Carnapian inductive logic's good Lakatosian marks might not have been caused by reappraisal-meriting factors but by a historical accident. Luckily for Carnapian inductive logic, it had a relatively short lifespan; during this time, I have argued, it did not develop any new hard-core commitments or diverge greatly from its original plans. As a result, we assessed Carnapian inductive logic as having had adequate heuristic power. However, if Carnapian inductive logic had had time to develop, it might not have proved so rigid. Even if it had, this might not have been such a good thing as Lakatos thought.

Contextual factors

Finally, there are a range of relevant contextual factors that are not obviously captured within Lakatos's framework. A number of famous critiques, which I discuss in chapter 3, argue that Carnapian inductive logic was fundamentally unviable. If any of these critiques were successful, then Carnapian inductive cannot have much to offer contemporary philosophy. Studying these critiques, and Carnapian inductive logicians' responses to them, is also important in order to gain a complete historical understanding.

Another important contextual factor is the existence of rival research programmes, particularly subjective and objective Bayesian inductive logic, that address similar problems to Carnap's. If Carnapian inductive logic cannot do anything that these programmes cannot do better, this would be another reason not to carry out a substantive reappraisal.

Defending the methodology of scientific research programmes against all of these objections would take us too far away from the main focus of this thesis. Rather than doing so, I will simply note that we would do best not to rely too heavily

on it. In the remaining chapters, I shall present independent reasons to be optimistic about the prospects of Carnapian inductive logic, taking into account the kind of contextual factors that the methodology of scientific research programmes arguably underplays. In chapter 3, I will address what I consider the most important critiques of Carnapian inductive logic, aruging that they can be overcome. Chapters 4 and 5 compare Carnapian inductive logic with two rival research programmes, namely subjective and objective Bayesian inductive logic. Finally in chapter 6 I argue that Carnapian inductive logic has useful applications. These arguments will hopefully outweigh in the reader's mind any misgivings about Lakatos's methodology.

Before doing so, however, it will be useful to address some textual issues relating to the reading that I presented above.

2.9 Additional interpretative points

The account of Carnapian inductive logic that I present above relies on a reading of Carnap's publications on inductive logic that differs somewhat from most others that appear in the secondary literature. My reading is substantially in agreement with those of several authors who collaborated with Carnap: see, for example, Jeffrey (1975), Jeffrey (1973) and Kemeny (1963b). However, unlike my account, these works tend simply to assert positions that Carnap held, rather than arguing for them on the basis of what Carnap wrote.

Before we proceed to the task of evaluating the critiques of Carnapian inductive logic, it will be useful to discuss some ways in which my account of Carnapian inductive logic differs from others. These differences are to some extent implicit in the discussion above, but here they are spelled out more fully.

2.9.1 Blank slate situations

In a late essay that appeared in several different versions—Carnap (1966a), Carnap (1968) and finally Carnap (1971b)—Carnap described a procedure for justifying inductive logical axioms involving what I shall call 'blank slate situations': scenarios where a certain kind of knowledge is completely unavailable. Carnap's remarks on blank slate situations have caused some misconceptions, which I aim to dispel in this section.

Carnap thought that inductive logical axioms could sometimes be justified on the grounds that they formalise rationality conditions that apply to agents who find themselves in blank slate situations. His reasons for focusing on blank slate rationality, rather than rationality in general, were as follows. Carnap claimed that, subject to certain epistemological assumptions, which he admitted were very unrealistic, the state of belief of an agent who starts in a blank slate situation and then receives some evidence can be identified with the state

of belief that, in the blank slate situation, they had thought that they *would* develop given this contingency. For this reason, Carnap believed that, in a certain sense, blank slate states of belief indicate permanent epistemic dispositions, whereas other states of belief reflect mere momentary opinions. Carnap thought that just as, when making ethical judgments, a person's permanent underlying character might be considered more pertinent than their motives at a particular moment, so an underlying epistemic disposition is more pertinent for assessing a person's rationality.

Unfortunately, given their prominence in Carnap's late work, his remarks on exactly what an agent can be taken to know in a blank slate situation were somewhat incomplete. What explanation there is comes largely from the following quotation. In it, Cr_0 is an 'initial credence' function representing the blank slate state of belief belonging either to a robot or a human being.

> How can we understand Cr_0? In terms of the robot, Cr_0 is the credence function that we initially build in and that he transforms step for step, with regard to the incoming data, into the later credence functions. In the case of a human being X, suppose that we find at the time T_n his credence function Cr_n. Then we can, under suitable conditions, reconstruct a sequence $E_1, \ldots, E_n$, the proposition K_n and a function Cr_0 such that:
> (a) $E_1, \ldots, E_n$ are possible observation data,
> (b) K_n is defined by $[K_n = \cap_n E_n]$,
> (c) Cr_0 satisfies all requirements of rationality for initial credence functions, and
> (d) the application of $[Cr_n(H) = Cr_0(H \mid K_n)]$ to the assumed function Cr_0 and K_n would lead to the ascertained function Cr_n.
>
> (Carnap, 1971b, p.18, square brackets added)

This passage does not specify in full exactly what the blank slate agent should be taken to know, other than implying that they have no 'observation data'. Carnap provides some clarification a few pages later at (Carnap, 1971b, p.23), stipulating that blank slate agents do not know about any differences between individuals. We are asked to consider two propositions H and H' which differ only in that H' attributes to one individual exactly the properties that H attributes to another.

> H and H' have exactly the same logical form; they differ merely by their reference to two distinct individuals. These individuals may happen to be quite different. But since their differences are not known to X at time T_0, they cannot have any influence on the Cr_0-values of H and H'....
>
> Suppose that X is a robot constructed by us. Since the propositions H and H' are alike in all their logical properties, it would be entirely arbitrary and therefore unreasonable for us to assign to

them different Cr_0 values in the construction of X.
(Carnap, 1971b, p.23)

These remarks seem to motivate imposing the axiom of constant exchangeability on all initial credence functions. This line of reasoning is in apparent tension with what Carnap goes on to say later in the same volume, when at (Carnap, 1971a, p. 120) he explicitly considers situations in which a rational agent's credences might not obey this axiom.

Carnap seems to be arguing that if two propositions share the same 'logical form', then an agent or robot in a blank slate situation should believe them to the same degree. Furthermore, H and H' in the example above are examples of propositions that share the same logical form. However, the situation is still not entirely clear: Carnap does not specify in general which propositions have the same logical form. What if, for example, two propositions differ only by reference to two distinct properties?

The Manchester reading

The Manchester school of inductive logicians pursue an approach that effectively fills this gap in Carnap's account of blank slate situations. This approach suggests the following 'Manchester reading' of Carnap's remarks on blank slate situations.[6]

The Manchester reading gives a fully fleshed out account of the blank slate situation: call this situation the Manchester blank slate situation. An agent in this blank slate situation is presented with an inductive logical domain D, but given absolutely no other information. In particular, they do not know how D should be interpreted, or what propositions, if any, its members represent. Nonetheless, the agent must choose a credence function representing their degrees of belief in whatever propositions the members of D happen to stand for. The Manchester reading of Carnap's remarks claims that the blank slate situation that he had in mind was the Manchester blank slate situation.

The Manchester reading coheres with Carnap's stipulation that propositions with the same 'logical form' should be given equal credence in blank slate situations, and might seem at first glance to chime with his remarks about the distinction between 'pure' and 'applied' inductive logic. On a natural account, pure inductive logic might be thought to be directed to the task of finding axioms that formalise rationality constraints for agents in the Manchester blank slate situation, whereas 'applied' inductive logic might be thought to focus on axioms that formalise rationality constraints for agents with background knowl-

---

[6]This choice of terminology is perhaps slightly misleading, as it suggests that views of members of the Manchester school about blank slate situations are motivated by a reading of Carnap's writing on this topic. In fact, the Manchester school's views predate their study of the relevant passage from Carnap (1971b).

edge, including about how to interpret the domains of their systems of inductive logic.

The follosing quotations give the gist of the Manchester reading. The reader should note that 'pure inductive logic' is here used in the liberal sense of the Manchester school, rather than in Carnap's more restricted sense.

> To capture the underlying problem that PIL [pure inductive logic] aims to address we can imagine an agent who inhabits some structure $M \in \mathcal{T}L$ but knows nothing about what is true in M. Then the problem is
>
> > $\mathcal{Q}$: In this situation of zero knowledge, logically, or rationally, what belief should our agent give to a sentence $\theta \in SL$ being true in M?
>
> There are several terms in this question which need explaining. Firstly zero knowledge means that the agent has no intended interpretation of the a_i nor the R_j.
> . . .
> In a way this is at the heart of the difference between the 'Pure Inductive Logic' proposed here as Mathematics and the 'Applied Inductive Logic' of Philosophy. For many philosophers would argue that in this latter the language is intended to carry with it an interpretation and that without it one is doing Pure Mathematics not Philosophy. It is the reason why Grue is a paradox in Philosophy and simply an invalid argument in Mathematics. Nevertheless, mathematicians or not, we all need to be on our guard against allowing interpretations to slip in subconsciously. Carnap himself was very well aware of this distinction, and the dangers presented by ignoring it, and spent some effort explaining it in [Carnap (1971a)]. Indeed in that paper he describes Inductive Logic as the study of the rational beliefs of just such a zero knowledge agent, a 'robot' as he terms it. (Paris and Vencovská, 2015, p. 10, square parentheses added)

> For Carnap's strategy to work, one must find a blank slate function, i.e., a probability function that captures the case in which we have zero information. One might naturally object that, in practice, we never really have zero information. In response, one may accept that, as a matter of fact, we always take a large variety of propositions for granted in our reasoning, but maintain that one can still think counterfactually about a blank slate function. Thus one can ask, what would be a reasonable probability function were one to have no information?
>
> This demands that one imagine a scenario in which there is a total absence of contextual and premiss propositions. In particular, there should be no contextual information about the meanings of the non-logical symbols in the language $\mathcal{L}$. Thus one must take the language to be uninterpreted: one must work in *pure* inductive logic.

(Williamson, 2015, p. 50)

I think that the Manchester reading does not quite capture what Carnap was getting at. First, the Manchester reading clashes with Carnap's remarks about the distinction between pure and applied inductive logic; second, it is at odds with some specific things he says about blank slate situations; finally, it suggests a view about logical languages that Carnap would not have agreed with.

Pure and applied inductive logic

Close attention to Carnap's remarks on the distinction between pure and applied inductive logic seems to undermine the Manchester reading. Carnap was clear that, according to his use of the term 'applied', arguments to do with the justification of inductive logical axioms should proceed only within applied inductive logic. Since Carnap introduced the blank slate situation in the context of an investigation into which inductive logical axioms are justified, his remarks must be read in the context of applied, rather than pure, inductive logic. This impression is reinforced by the following quotation, in which Carnap explicitly puts questions of rationality and credibility in the 'applied' pigeonhole.

> It is the task of pure inductive logic to state axioms for c-functions and derive theorems from these axioms. In applied IL, the theorems are used for practical purposes, e.g., for the determination of the credibility of a hypothesis under consideration in a given knowledge situation, or for the choice of a rational decision.
> (Carnap, 1971a, p.105)

For Carnap, pure inductive logic could only state axioms and derive theorems from them. It was only within applied inductive logic that questions about whether particular axioms formalise rationality constraints could arise. This would be the case even for rationality constraints that are so general that they apply to agents who find themselves in the Manchester blank slate situation.

It is informative to recall the analogy that Carnap drew between pure and applied inductive logic and pure and physical geometry. Just as, Carnap thought, the question of which structures represent physical space does not arise within pure geometry, the question of which systems of inductive logic represent rationality constraints did not arise within pure inductive logic.

Carnap's remarks on pure and applied inductive logic therefore do not support the Manchester reading. Moreover, without the motivation of working within pure inductive logic, there is little reason to think that Carnap would have had such a stark scenario as the Manchester blank slate situation in mind. The situation where an agent has to attach credences to an uninterpreted domain is very atypical of ordinary inductive reasoning, which was Carnap's main focus.

Carnap's remarks on blank slate situations

Some aspects of Carnap's remarks on blank slate situations give further reasons

to be dubious that he had the Manchester blank slate situation in mind.

First, it is clear that Carnap envisaged some non-logical background knowledge being available in blank slate situations. In both the case of the robot and the idealised human Carnap talks about 'observations'. If either agent were meant, as in the Manchester blank slate situation, to have access to no knowledge besides possession of an uninterpreted inductive logical domain, this term would be inapt: the agent would not be able to tell whether a given domain-member should be thought of as representing an observation. This would have been odd if Carnap had had the Manchester blank slate situation in mind.

Another reason comes from Carnap's comments on the axiom of constant exchangeability, which his blank slate scenario is meant to help justify. Carnap thought that constant exchangeability formalises a rationality constraint in his blank slate situation, but also wrote that it should not always be adopted. In particular, Carnap thought that constant exchangeability is typically inappropriate in the case of 'coordinate languages'.

> It is to be noted that symmetry with respect to individuals should be required only for those languages in which all individual constants have the same logical nature. The situation is different for a *coordinate language*, i.e., one in which the standard individual expressions indicate the positions of the individuals...
> (Carnap, 1971a, p. 119, emphasis original)

Whether a given language L is a coordinate language does not depend on its formal properties, but rather on a question of interpretation. Specifically, it must be ascertained whether certain elements are taken to represent positions of individuals that affect the 'logical nature' of the constants that they attach to. Suppose that L is a predicate language with constants $a_1, a_2, \ldots$, where the indices $1, 2, \ldots$ are taken to represent inductively relevant positions - for example the position of an experiment in a sequence where the result of one round is relevant to the result of the next. Then the constants $a_1, a_2, \ldots$ do not have 'the same logical nature', and L is a coordinate language. On the other hand, if the indices were taken to represent inductively irrelevant positions in a sequence of coin flips, then L would not be a coordinate language.

Suppose that an agent in the Manchester blank slate situation is presented with a domain SL. Depending on how L's indices are interpreted, L may be a non-coordinate language, in which case constant exchangeability might be appropriate, or it may be a coordinate language, in which case constant exchangeability would be inappropriate. The agent cannot tell which of these possibilities is the case because, by stipulation, they do not know how to interpret L.

There therefore seem to be little grounds, according to Carnap's own criteria, to base an argument for constant exchangeability on the Manchester blank slate situation. Yet Carnap clearly *did* want to use his blank slate situation to motivate constant exchangeability. The best explanation for this state of affairs is

that the blank slate situation that Carnap had in mind was not the Manchester one.

Voluntarism

A final, broader, reason to doubt that Carnap had the Manchester blank slate situation in mind has to do with his attitudes towards language. One of the central features of Carnap's mature philosophy was his emphasis on the freedom that human agents have to choose between different languages, and his rejection of the view that any particular language is unavoidable. He devotes a chapter of his *Introduction to the Philosophy of Physics* to arguing against what he calls the 'magical' view of language, according to which "there is some mysterious magical connection of some sort between certain words...and their meanings" (Carnap, 1966b, p. 115-116).

However, an agent in the Manchester blank slate situation finds themself in exactly the position that Carnap was so keen to stress was not the typical one. They are simply presented with a domain D and told that it represents something, but not told what that something is. They can choose their credibility function but not their language. Effectively, the agent is forced to adopt a very general form of the magical view of language.

While this link does not show definitively that Carnap could not have had the Manchester blank slate situation in mind, the fact that this scenario is so incongruous with his general attitude weighs against the Manchester reading.

An alternative reading

Here is an alternative reading of how Carnap intended his blank slate situation to be understood. An agent in a 'Carnap blank slate situation', I believe, should be thought of as a scientist who is about to conduct a series of experiments which they wish to learn from. They have no *observational* information in the narrow sense that they have not yet learned the results of the experiments. However they should also be thought of as having some non-observational background knowledge - in particular they know what the possible results of the experiments are and know enough about the experiments to judge that their order conveys no information.

This reading agrees with what Carnap says about the blank slate situation. Knowing the nature of the setup, the agent is in a position to treat the outcomes of the experiments as "incoming data" and determine a strategy for learning from them. When Carnap introduces the blank slate situation in (Carnap, 1971b, Ch. 4) he talks only about the agent lacking 'observational knowledge'; the fact that other kinds of knowledge are not discussed supports my reading's attribution of non-observational information to the blank slate agent. Finally, there is a clear sense in which two propositions that are the same except for referring to different experiments have the same "logical form": the agent's

information about the two propositions is the same, given the assumption that the order of experiments conveys no information. This explains why Carnap used his blank slate situation to motivate constant exchangeability.

Conclusion

Carnap's broad aim in Carnap (1971b) was to show that his theory of inductive logic could be connected with the then-standard theory of rational decision-making under uncertainty. In his writing about inductive logic in general, there is little discussion of the blank slate situation: it provides just one out of many possible methods for showing that a system of inductive logic usefully formalises an inductive assumption. For this reason I do not think blank slate situations are an essential component of Carnapian inductive logic and so did not discuss it in the section above.

Nonetheless, it is important to note that the Manchester blank slate situation, according to which an agent is presented with an uninterpreted inductive logical domain, seems not to be what Carnap had in mind, and can be replaced by an alternative account.

2.9.2   The term 'logic'

In the context of philosophy, the term 'logic' typically connotes formal rules concerning deductive entailment. Given this context, using the term 'logic' in connection with inductive reasoning might be considered inappropriate. Leitgeb, for example, argues that doing so is uninformative and that 'logic proper' encompasses only theories of deductive consequence relations:

> In [a] sense, the probabilistic explication of the confirmation of hypotheses (as initiated by Carnap 1950) may, for example, be regarded as a kind of probabilistic (or inductive) logic... On the other hand, if used in such a broad manner, the label 'probabilistic logic' is no longer particularly informative as far as its 'logic' component is concerned. In this chapter, we will restrict the term 'logic' to logic *proper*: a logic or logical system is a triple of the form $\langle \mathcal{L}, \vDash, \vdash \rangle$ where (i) $\mathcal{L}$ is a formal language, (ii) $\vDash$ is a semantically (model-theoretically) specified relation of logical consequence defined for the members of $\mathcal{L}$ (iii) $\vdash$ is a proof-theoretically (in terms of axioms and rules) specified relation of deductive consequence for the members of $\mathcal{L}$.
>
> (Leitgeb, ming, p. 1-2, square brackets added)

Howson puts a related point somewhat more bluntly:

> Talk of probabilities on sentences of a language brings us naturally, and indeed inevitably, to the case of Carnap. In [*Logical*

> *Foundations of Probability*] he inaugurated a research tradition of
> combining logic with probability in the following way: define a real-
> valued, finitely additive normalised function on the sentences of a
> formal language. ...Carnap called the result logical probability,
> without much in the way of argument to support this terminology,
> and none of it very convincing. If one wanted to be funny one might
> describe Carnaps model as affirming something like
>
> $$\text{logic} + \text{probability} = \text{logical probability}$$
> $$\approx \text{cheese} + \text{soufflé} = \text{cheese soufflé}$$
>
> (Howson, 2009, p.1)

The question of whether Carnapian inductive logic should really be called 'logic'
is fundamentally terminological, and therefore not critical to the success or
failure of Carnap's research programme. Nonetheless, it is interesting to consider
why Carnap chose to use the term 'inductive logic'. Contrary to Howson, Carnap
did present reasons for this choice. First, systems of inductive logic are logical
entities according to Carnap's conception of logic. Second, Carnap thought that
there were parallels between his research programme of inductive logic and that
of deductive logic.

Carnap's conception of logic

According to Carnap's conception, logic comprises one component of the gram-
mar of a language, indicating which sentences follow from which others. He
expressed this view as follows:

> In [the logician's] sense, a language is constituted by rules for
> a vocabulary, rules for building sentences, rules for making logical
> deductions from those sentences, and other rules.
> (Carnap, 1966b, p. 58-59)

On this conception, choices between logical systems, just like other aspects of
language-choice, are essentially matters of convention. Carnap is famous for
having advocated tolerance between logical systems for this reason. He put this
view picturesquely as follows:

> In logic, there are no morals. Everyone is at liberty to build his
> own logic, i.e. his own form of language, as he wishes. All that is
> required of him is that, if he wishes to discuss it, he must state his
> methods clearly, and give syntactical rules instead of philosophical
> arguments.
> (Carnap, 1937, p.52)

While Carnap later came to acknowledge that semantic rules as well as syntactic
ones are indispensable, this remained his view during his later work. This can
be seen from his remarks at the end of Carnap (1950a):

> To decree dogmatic prohibitions of certain linguistic forms in-
> stead of testing them by their success or failure in practical use, is
> worse than futile; it is positively harmful because it may obstruct
> scientific progress. ...
> Let us be cautious in making assertions and critical in examining
> them, but tolerant in permitting linguistic forms.
> (Carnap, 1950a, § 5)

According to Carnap's conception of logic, then, any system of rules specifying what follows from what within a language merits being called a logic. Thus Carnap defended his use of the term 'inductive logic' by pointing out that his programme aimed to specify formal rules according to which certain statements about measure functions follow from certain other ones:

> All axioms of inductive logic state relations among values of m
> or c as dependent only upon the logical properties and relations of
> the propositions involved (with respect to language systems with
> specified logical rules). Inductive logic is the theory based upon
> these axioms. It may be regarded as part of logic in view of the fact
> that the concepts occurring are logical concepts.
> (Carnap, 1971b, p. 26)

Analogies with deductive logic

Carnap's second reason for using the term 'logic' in connection with inductive reasoning was that there are similarities between inductive and deductive logic. At (Carnap, 1950b, § 43) he explains at length what he thought the similarities were. The guiding thought behind this analogy can be illustrated as follows.

Carnapian inductive logic is concerned primarily with formal axioms that pick out measure functions representing degrees of plausibility that agree with given assumptions about how inductive reasoning should proceed. Deductive logical semantics can be thought of as concerned with formal admissibility conditions that pick out valuation functions representing configurations of truth or falsity of propositions that agree with given assumptions. Looked at in this way, inductive logical axioms are analogous to admissibility constraints on valuation functions. Whereas admissibility constraints represent intuitions about truth— one might label them 'deductive assumptions'—axioms in Carnapian inductive logic represent inductive assumptions about plausibility.

To take an example, it is normal in classical logic to require valuation functions to behave in such a way that $v(\theta \wedge \phi) \leq min(v(\theta), v(\phi))$. This is because, according to a commonly made deductive assumption, it is not possible for a conjunction to be more true than either of its conjuncts. Similarly, a useful inductive logical axiom stipulates that $m(\theta \wedge \phi) \leq min(m(\theta), m(\phi))$. This constraint is useful because according to a commonly made inductive assumption, a conjunction cannot be more plausible than one of its conjuncts.

While not all admissibility conditions for valuation functions have such neat inductive logical analogues, there are clearly structural similarities between the formal frameworks that describe deductive and inductive assumptions. Arguably, if the term 'logic' is merited in one case, then it should also be used in the other.

Conclusion

It strikes me that Carnap's conventionalist conception of logic was sensible, and that there are indeed strong analogies between Carnapian inductive logic and deductive logic. However, as noted in the introduction, I do not intend to defend his use of the term 'logic' in this thesis. It is sufficient for my purposes to make clear that Carnap did have principled reasons for using the term.

2.9.3 Probability$_1$

In Carnap (1950b), Carnap uses the term 'probability$_1$' to describe the collection of inductive assumptions that collectively constitute the norms of inductive reasoning in science and everyday life, and which he aims to explicate in that book. Since I think that Carnap's use of term is somewhat loaded—it implies a certain degree of agreement about rules of inductive reasoning, and a special role for probability functions—I did not mention it in the account sketched above.

Nonetheless it is useful to discuss Carnap's remarks about probability$_1$ carefully, and in particular his assertions that probability$_1$ is 'objective', 'logical' and 'analytic', as these have been widely misconstrued.

Carnap explains that 'probability$_1$' is the concept of "degree of confirmation"(Carnap, 1950b, p.19) or "weight of evidence"(Carnap, 1950b, p.163) which is in " general"(Carnap, 1950b, p.19) and "presystematic"(Carnap, 1950b, p.23) use, for example by scientists trying to work out to what extent observational evidence support particular hypotheses (Carnap, 1950b, p.20).

Carnap emphasises that that probability$_1$ is a 'logical' concept, putting the point as follows.

> We call [the question to what degree one statement confirms another]
> a logical question because. . . the problem whether and how much h
> is confirmed by e is to be answered merely by a logical analysis of h
> and e and their relations. . . . [O]nce h and e are given, the question
> mentioned requires only that we be able to understand them, that is,
> to understand their meanings and to grasp certain relations which
> are based upon their meanings. (Carnap, 1950b, p. 20)

It is important to note that, despite the use of the symbols 'h' and 'e' as abbreviations for longer statements, and the reference to 'logical analysis', Carnap

is here making a socio-linguistic claim about how inductive assumptions are expressed in ordinary non-symbolic discourse. If it was not a part of ordinary discourse, then probability$_1$ would not, as Carnap claims, be in "general" and "presystematic" use. Carnap's claim is that, when someone utters a sentence such as 'Mrs X's testimony strongly suggests that the defendant is guilty', they do not thereby express a claim about the testimony, the defendant, or any other contingent matter. Instad, Carnap thought, this kind of utterance expresses a certain meta-linguistic about the semantic rules of the prevailing language, namely that the sentence expressing the testimony is related in a certain way to the sentence expressing the defendant's guilt. Whether or not this relationship obtains is a matter to be determined not by investing the world but by through logical analysis.

The same might be said about the phrase 'Either it is true that it is raining or it is false that it is raining'. This phrase, Carnap would say, does not express a fact about the weather, but rather a fact about the semantic rules of its language, namely that its truth predicate obeys the law of excluded middle.

All kinds of inductive assumptions, and not just the generally used ones that constitute probability$_1$, can be thought of as semantic rules in this way, and are 'logical' in this sense that Carnap attributed to probability$_1$.

The second fact about probability$_1$ that Carnap emphasises is that, in virtue of being logical, it is an 'objective' concept. Logical concepts are 'objective', Carnap says, because

> a sentence which ascribes one of these concepts in a concrete case. . . is complete without any reference to the properties or behaviour of a particular person.
> (Carnap, 1950b, p.38)

In contrast, Carnap points out, the concept of 'familiarity' is subjective because, grammatically, it can only correctly be ascribed if it is clear *to whom* something or other is familiar. Carnap's assertion that probability$_1$ is 'objective' is therefore very weak: all inductive assumptions, including ones that sanction very bad inductive reasoning, are objective in exactly the same sense.

Finally, Carnap sometimes says that true statements about probability$_1$ are analytic—that is, true in virtue of meaning rather than in virtue of correspondence with the world. This is a trivial consequence of Carnap's view that probability$_1$ is an essentially a semantic phenomenon. Statements about probability$_1$ were, according to Carnap, typically made in the context of an implicitly accepted system of semantic rules: if true at all, such statements must be consequences of these rules and, in this sense, true in virtue of meaning.

Thus, Carnap's assertions that probability$_1$ is a 'logical', 'objective' and 'analytic' concept does not single it out in comparison with other collections of inductive assumptions, or say anything at all about his systems of inductive logic. Rather, they merely serve to emphasise Carnap's claim, which is perhaps

not entirely uncontroversial, that real-world inductive assumptions are essentially expressions of semantic rules. While this claim is somewhat difficult to appraise, as the logico-grammatical structure of natural language is difficult to pin down, it is not important for the evaluation of Carnapian inductive logic. All that is really required is that real-world inductive assumptions, whatever their true nature, are germane to the kind of formalisation that Carnapian inductive logic proposes.

Chapter 3

Critiques

3.1 Introduction

This chapter defends Carnapian inductive logic against some of the most prominent attempts to show that it was a bad research programme.

I discuss at length what I consider to be the most important critiques. The first was made by Lakatos, who argued that Carnapian inductive logic failed due to being unable to identify true axioms. I then address the 'grue' critique, which alleges that Carnapian inductive logic is undesirably 'syntactic', and a critique due to Salmon based on logical entailment. Finally I consider three additional critiques that arise from worries about symmetry, general statements and long-run reliability.

3.2 Lakatos's critique

As well as supplying a useful framework with which to analyse Carnapian inductive logic, Lakatos was also one of its most influential critics. His highly-cited article Lakatos (1968a) claims that, as it developed, Carnapian inductive logic became degenerate by disobeying its negative heuristic - the most serious way in which a research programme can fail according to the methodology of scientific research programmes. According to Lakatos, Carnapian inductive logic gave up a hard-core commitment to searching for true inductive logical axioms in order to pursue a less ambitious goal. This account of how Carnapian inductive logic developed is still widely believed, and although modern philosophers rarely explicitly apply the methodology of scientific research programmes to Carnapian inductive logic, I believe that many would agree with Lakatos's assessment.

However, Lakatos was mistaken: Carnapian inductive logic never made any commitment, hard-core or not, to the existence of true inductive logical axioms. In fact, it was a central feature of Carnap's mature philosophy that logical axioms should not be thought of as true or false, but only as more or less useful.

In this section I reproduce Lakatos's argument, argue that it was faulty and attempt to explain what might have caused Lakatos's error.

3.2.1   Historical background

In July 1965, a section of the International Colloquium in the Philosophy of Science in London was devoted to "The Problem of Inductive Logic". The conference, organised by Imre Lakatos and Karl Popper among others, was attended by Rudolf Carnap, and focused on Carnapian inductive logic. Three years later, a proceedings volume entitled *The Problem of Inductive logic* was published, edited by Lakatos. The volume consists of a series of what Lakatos called "rational reconstructions"[1] of the conference talks and ensuing discussions, in which each speaker submitted an essay reflecting their talk, followed by critical responses from the discussants and then rejoinders from the original speaker. This series is followed by Lakatos (1968a), a long essay originally conceived as a response to Carnap's contribution, Carnap (1968)[2].

Unlike the other essays in the proceedings volume, Lakatos (1968a) does not attempt to reconstruct the discussions from the London conference. It is only loosely related to Carnap (1968) and is not followed by a final response from Carnap, who did not see the final version prior to publication[3]. Instead, Lakatos attempts to set contemporary discussions about inductive logic in the context of a dialectical history of thought culminating in a conflict between a tendency represented by Carnapian inductive logic—'neo-classical empiricism'—and 'critical empiricism', a position embodied by Popper and himself. As part of this presentation Lakatos argued that Carnapian inductive logic had made what he calls "a degenerating problem-shift".

With this locution, Lakatos was implicitly invoking his methodology of scientific research programmes, which he was working on simultaneously[4]. At (Lakatos, 1968a, p. 357), Lakatos refers to "Carnap's research programme of inductive logic", showing that he was thinking about it in these terms.

Carnap did not respond to Lakatos's essay in any published work, though he did send Lakatos a letter Carnap (1967a) in response to an early draft Lakatos (1967c). In the letter Carnap made several critical remarks and asked Lakatos not to publish the essay in the proceedings volume. Unfortunately the letter does

---

[1]Lakatos describes the contributions as such in Lakatos (1966).

[2]Lakatos explains how he came to write his essay in Lakatos (1967b).

[3]Lakatos points this out at (Lakatos, 1968a, p.1).

[4]The first major statement of Lakatos's methodology, Lakatos (1968b), was published only a few months after Lakatos (1968a).

not reveal Carnap's opinion about Lakatos's main argument for the degeneracy of his research programme, which did not appear in the draft that he received.

According to Lakatos's argument, the primary goal of Carnapian inductive logic was to identify a formal system that was in some sense true, and could be used to pronounce on the rationality or irrationality of beliefs and scientific theories. However, Lakatos's account continued, this task proved insurmountable, forcing Carnapian inductive logic to become a 'mere calculus of coherent beliefs'. Abandoning a core commitment in this way, Lakatos concluded, implicitly invoking his nascent methodology, was a clear sign that Carnapian inductive logic was a degenerate research programme.

Lakatos's argument was well read and cited: at the time of writing Google scholar records 285 citations of Lakatos (1968a). Today his view that Carnapian inductive logic was degenerate is widely accepted, though Lakatos did not achieve his goal of directing support towards the Popperian programme of 'critical empiricism' that he advocates in the second section of Lakatos (1968a).

3.2.2 Lakatos's argument

Lakatos's argument for Carnapian inductive logic's degeneracy begins by offering an account of the motivating ideas and goals behind Carnap's research programme:

> Thus Carnap—following the Cambridge school (Johnson, Broad, Keynes, Nicod, Ramsey, Jeffreys), Reichenbach, and others—set out to solve the following problems: (1) to justify his claim that the degree of confirmation satisfies Kolmogorov's axioms of probability, (2) to find and justify further secondary adequacy requirements for the determination of the sought-for measure function; (3) to construct—piecemeal—a complete, perfect language of science in which all propositions of science can be expressed; and (4) to offer a definition of a measure function which would satisfy the conditions laid down in (1) and (2).

> Carnap thought that while science was conjectural, the theory of probabilistic confirmation would be *a priori* and infallible: the axioms, whether primary or secondary, would be seen to be true in the light of *inductive intuition* and the language (the third ingredient) would of course be irrefutable, for how can one refute a language? (Lakatos, 1968a, p.323-324, emphasis original)

Lakatos thought that early Carnapian inductive logic aimed to find a system of inductive logic whose axioms were 'true in the light of inductive intuition' and whose domain was the set of sentences of a 'complete', 'perfect', 'irrefutable' language. The resulting system would be 'a priori and infallible'. It is safe to conclude, therefore, that Lakatos thought that the aim of finding inductive

logical axioms that are true in some sense was part of the hard core of early Carnapian inductive logic.

However, according to Lakatos, careful attention to scientific reasoning shows that it is impossible in principle to identify such a system of inductive logic. The language of science, he argues at (Lakatos, 1968a, § 2.3(a)), has sometimes changed so radically as to alter which propositions are thought to provide inductive support to which others. Since such volatility is likely to continue, Lakatos concludes, it will never be possible to know for sure if any true axioms have been found:

> The growth of science may destroy any particular confirmation theory: the inductive machine may have to be reprogrammed with each new major theoretical advance.
> (Lakatos, 1968a, p.364)

Lakatos claims that Carnapian inductive logic failed to anticipate this difficulty. When the difficulty was encountered, he says, Carnapian inductive logic was forced to relinquish its original goal of identifying true axioms. Rather than acting as an "inductive judge", ruling on substantive questions of rationality, he says, late Carnapian inductive logic sought to achieve only the less ambitious task of formalising inductive assumptions.

> The abdication of the inductive judge is complete. He promised to hand down judgement on the rationality of beliefs; now he is ending up by trying to supply a calculus of coherent beliefs on whose rationality he cannot pronounce.
> (Lakatos, 1968a, p.372)

Lakatos concludes that this change, whereby a former component of Carnapian inductive logic's hard core was discarded, radically disobeyed its negative heuristic. Carnapian inductive logic therefore failed the test of heuristic power and became a degenerate research programme:

> A mere calculus of coherent beliefs can at best have marginal significance relative to the central problem of the philosophy of science. Thus, in the course of the evolution of the research programme of inductive logic its problem has become much less interesting than the original one: the historian of thought may have to record a 'degenerating problem-shift'.
> (Lakatos, 1968a, p.373)

One might think that Lakatos could not have been applying his methodology to Carnapian inductive logic because the latter is a philosophical, rather than a scientific, research programme. In fact, Lakatos was prepared to apply his methodology to to philosophical research programmes: at (Lakatos, 1978a, p.92) he says that his method "can be generalised to any sort of rational discussion and thus serve as tools for a general theory of criticism", referring to Carnapian inductive logic explicitly as a "degenerating programme".

3.2.3 Why Lakatos was mistaken

Lakatos's main error was to assume that early Carnapian inductive logic aimed to true axioms. In fact, as we saw in the previous chapter, and particularly section 2.4.3, early Carnapian inductive logic was committed to the view that axioms, like other parts of systems of inductive logic, were impossible to justify on theoretical grounds, but could only be justified pragmatically as more or less apt for particular applications.

In other words, Carnapian inductive logic always aimed simply to "supply a calculus" with which patterns of inductive reasoning can be formalised. It never aimed to find axioms that were true in any sense, let alone 'a priori', 'infallible', etc. Nor did it aim to pass judgement on the rationality of scientific theories, except insofar as such judgement can derive from formal statements of principles which had previously been articulated unclearly or inconsistently.

The fact that Carnapian inductive logic never aimed to act as an "inductive judge" can be seen from the following two quotations from Carnap's early work:

> The system of inductive logic here proposed. . . is intended as a reconstruction. . . of inductive thinking as customarily applied in everyday life and science. However, it is not meant merely as an uncritical representation of customary ways of thinking with all their defects and inconsistencies, but rather as a rational, critically corrected reconstruction. (Carnap, 1950b, p.576)

> First I wish to emphasize that inductive logic does not propose new ways of thinking, but merely to explicate old ways. It tries to make explicit certain forms of reasoning which implicitly or instinctively have always been applied both in every day life and in science. This is analogous to the situation at the beginning of deductive logic. Aristotle did not invent deductive reasoning; that had gone on as long ago as there was human language. If somebody had said to Aristotle: "What good is your new theory to us? We have done well enough without it. Why should we change our ways of thinking and accept your new invention?", he might have answered: "I do not propose new ways of thinking, I merely want to help you to do consciously and hence with greater clarity and safety from pitfalls what you have always done. I merely want to replace common sense with exact rules" It is the same with inductive logic.
> (Carnap, 1953, p.189)

Carnap was not aiming to demonstrate anything that was not already clear to everyday inductive reasoners, but only, just as in Aristotle's case, to articulate clearly the standards that they already use. In particular, just as Aristotle's deductive logic was not designed to give reasons to choose one assumption about how deductive reasoning ought to proceed over another, so inductive logic was not designed to show that particular inductive assumptions are justified.

It might be thought that, even if Lakatos was mistaken about Carnapian inductive logic wanting to pass judgement on the question of which inductive assumptions are rational, he was still right that systems of inductive logic were meant to be true in a narrower sense, namely that they were meant to express truths about generally accepted inductive assumptions. Systems of inductive logic, were certainly meant to be 'similar' to the inductive assumptions they formalise: we have seen in section 2.4.1 that this was one of the features that could make an explication successful. However, the situation is not straightforward because similarity to the explicandum is only one of many possible virtues—it is possible that a system of inductive logic that only very faintly resembled any actually instantiated assumptions could still usefully explicate them. In any case, even if Carnapian inductive logic did aim at truth in this sense, this aim did not change from early to late Carnapian inductive logic. As a result, Lakatos's argument for degeneracy would still fail to go through.

3.2.4   Why did Lakatos go wrong?

The reader might wonder why Lakatos mis-characterised Carnapian inductive logic in this way. I think that the cause was a combination of some misleading writing in Carnap (1950b) as well as a tendency on Lakatos's part to focus on familiar problems.

Carnap's misleading writing

Lakatos's mis-characterisation may also have been caused by some unintentionally misleading writing in Carnap (1950b). In section 2.9.3 on probability$_1$ we saw that Carnap's discussion of this topic in Carnap (1950b) is best understood as consequences of Carnap's view that real-world inductive assumptions are essentially semantic. In particular, Carnap's claims that probability$_1$ is 'analytic', 'logical' and 'objective', were not as bold as they might seem at first glance.

Lakatos seems to have been misled by these remarks, as can be seen from his comments about analyticity and the similarity of Carnap and Keynes's views on probability.

Analyticity
Carnap often writes in Carnap (1950b) and elsewhere that statements about probability$_1$, as well as certain statements of inductive logic, can only be analytically, as opposed to synthetically, true or false. Correctly interpreted, these remarks refer only to 'internal questions' that concern the properties a given linguistic system: either a system of inductive logic or whichever semantic rules are expressed by natural-language inductive assumptions. Such statements, if true, relate only to the properties of formal systems and would therefore, at least according to Carnap, be uncontroversially analytic. That Carnap took

this view can be seen from his discussion of the contrast between logical and methodological problems at (Carnap, 1950b, § 44 A).

However, it is easy to misinterpret Carnap's remarks about analyticity as claiming that answers to external questions of inductive logic—those that concern choices between systems—or questions about how particular systems of inductive logic are related to the world, are also analytic, and thereby conclude that Carnap intended to find true axioms. Lakatos seems to have made precisely this mis-interpretation:

> ...Carnap took his inductive logic to be *analytic*.
> (Lakatos, 1968a, p. 324)

> Now Carnap's 'analytical' inductive principles consist partly of his explicit axioms, partly of his implicit meta-axioms...
> (Lakatos, 1968a, p.368)

Discussion of Keynes

Carnap discusses Keynes (1921) at (Carnap, 1950b, § 12A). Carnap notes that Keynes shared his view that probability should be thought of as in principle independent of particular agents, that is, as degree of *plausibility* rather than necessarily degree of *belief*. On this basis Carnap expresses substantial agreement with Keynes's project, even going so far as to note:

> ...the objective logical concept meant by Keynes is the same as what we call probability$_1$
> (Carnap, 1950b, p.44)

However, on the question of whether the rules that constitute probability$_1$ are true, Carnap took exactly the opposite view to Keynes. As can be seen from the following quotation, Keynes believed that probability relations between propositions existed in some sense, and that the task of inductive logic was to make true claims about these relations:

> We believe that there is some **real** objective relation between Darwin's evidence and his conclusions ...

> We are claiming, in fact, to cognise **correctly** a logical connection between one set of propositions which we call our evidence and which we suppose ourselves to know, and another set which we call our conclusions.... (Keynes, 1921, p. 5, emphasis added)

Carnap, on the other hand, did not think that probability$_1$, or any other collection of semantic rules, existed or could be identified:

> I hope that nobody will misinterpret my statement of the objectivity of logical relations as a metaphysical statement of the 'subsistence' of these relations in a Platonic heaven... (Carnap, 1950b, p.38)

At (Carnap, 1950b, p.45), Carnap explicitly distances himself from Keynes's view that it is possible to identify true probability relations using 'inductive intuition'.[5]

Lakatos seems not to have picked up on this distinction betwenn Carnap and Keynes's. He writes at (Lakatos, 1968a, p.323) that Carnap was "following the Cambridge school", which included Keynes, without pointing out the important differences between their views.

Lakatos's preoccupations

Lakatos's mis-characterisation of early Carnapian inductive logic may also have stemmed from a mistaken assumption that Carnap shared some of his own philo-sophical preoccupations. Lakatos was passionately interested in distinguishing science from pseudo-science, explaining why the history of science unfolds as it does and in putting this unfolding on an objectively rational basis. He seems to have thought that Carnapian inductive logic was directed at these problems too.

Lakatos suggested in a lecture delivered some years after the publication of Lakatos (1968a), that inductive logic—presumably including Carnapian induc-tive logic—aimed to solve the problem of distinguishing science from pseudo-science:

> If all scientific theories are equally unprovable, what distinguishes scientific knowledge from ignorance, science from pseudo-science?
> One answer to this question was provided in the twentieth cen-tury by 'inductive logicians'. Inductive logic set out to define the probabilities of different theories according to the available total evi-dence. If the mathematical probability of a theory is high, it qualifies as scientific; if it is low or even zero, it is not scientific. Thus the hallmark of scientific honesty would be never to say anything that is not at least highly probable.
> (Lakatos, 1978c, p.3)

Only probabilities that follow from true axioms would be useful for the task of distinguishing scientific theories from pseudo-scientific ones. If there were no facts about which theories were highly probable, there would be no grounds on which to distinguish pseudo-scientific theories.

Lakatos mistakenly thought that Carnapian inductive logic had, at least initially, sought to formalise scientific theories within a "complete language of science" (Lakatos, 1968a, p. 323) and to justify scientists' choices of theories.[6] Thus

---

[5]Carnap's acknowledgement in Carnap (1968) of the role of inductive intuition in inductive logic in no way represents a departure from this view. According to this paper, inductive intuition does not identify true systems of inductive logic: rather inductive logic seeks to codify our inductive intuitions.

[6]See (Carnap, 1950b, § 49) for early Carnap's explicit disavowal of these goals.

Lakatos seems to have thought that Carnapian inductive logic was interested in demonstrating the supposed objectivity of scientific knowledge.

This mistaken assumption would have made it more plausible that Carnapian inductive logic would have aimed to find true axioms. Indeed, unless it could identify such axioms, it is difficult to see how Carnapian inductive logic could be relevant to the question of whether scientific knowledge is objectively justified.

3.2.5   Summary

To sum up, Lakatos's argument for the degeneracy of Carnapian inductive logic depended on his claim that, in its early phase, Carnap's research programme sought to find true axioms. In fact, Lakatos was mistaken: Carnapian inductive logic only ever aimed to find pragmatically justified axioms. As a result, Lakatos's critique of Carnapian inductive logic was unsuccessful.

3.3   Grue, projectability and instantial relevance

Besides Lakatos's criticism, the most significant objection to Carnapian inductive logic has been the suggestion that it fails to deal satisfactorily with the 'grue' problem. In essence, this objection criticises Carnapian inductive logic on the grounds that it cannot formalise assumptions about both positive and negative instantial relevance, or, in Goodman's terms, 'projectability' and 'non-projectability'. Furthermore, Carnapian inductive logic gives no guidance as to when different kinds of assumption about instantial relevance are justified.

In this section I argue that Carnapian inductive logic *can* formalise assumptions of non-projectability just as well as it can formalise assumptions about projectability. On the other hand, it cannot show in what circumstances such assumptions are justified, but this is a very difficult task which Carnapian inductive logic never attempted and which no other research programme has ever achieved. Objections to Carnapian inductive logic based on 'grue' are therefore misplaced.

3.3.1   Historical background

The 'grue' problem emerged from an exchange between Carnap and Goodman starting with Goodman (1946), an essay in which Goodman raised some general objections to Hempel's qualitative theory of confirmation and Carnap's quantitative theory, which was derived from his work on inductive logic.

Carnap addressed Goodman's objection in Carnap (1947b), Goodman responded in Goodman (1947) and Carnap responded in Carnap (1948). Goodman drew

on this exchange to offer the definitive formulation of his objection in Goodman (1954), introducing the term 'grue'.

Many articles by other authors on 'the new riddle of induction' followed, including Jeffrey (1966), which I quote from below. Carnap later addressed Goodman's objection again in (Carnap, 1971a, § 4B).

Today it is widely believed that the grue problem presented a serious obstacle which Carnapian inductive logic failed to overcome.

3.3.2   The objection

The grue problem, it is generally agreed, has crucially to do with the distinction between 'projectable' and 'non-projectable' properties. While there has been some philosophical discussion about exactly how 'projectability' should be defined—see Earman (1989) for example—I shall here follow the definition that was implicitly agreed in the original discussion between Carnap and Goodman.

During the exchange with Goodman, Carnap defined projectability as follows:

> We call a property inductively projectible if the following is always the case: the higher the relative frequency of W in an observed sample, the higher is, on this evidence, the probability that a non-observed individual has the property W.
> (Carnap, 1947b, p.146)

Since Goodman did not dispute Carnap's definition, let us follow Carnap and define a property as 'projectable' in a given context if, according to inductive common sense, observed instances of that property make future instances more plausible in that context. Otherwise it is non-projectable. In other words, recalling the definition of the axiom of positive instantial relevance from section 2.4.2, a property is projectable whenever inductive common sense suggests an assumption of positive instantial relevance with respect to it.

As it is normally formulated, the grue problem asks us to consider a situation in which a series of emeralds are being examined. To fix a concrete image, imagine that emeralds are extracted from an underground mine by a kind digger who places them at regular intervals on a conveyor belt. When the emeralds reach the surface they are examined, one by one, by an emerald inspector with a penchant for inductive reasoning.

The property of 'grueness' is defined as follows: an emerald is 'grue' if it is either examined before a certain time t and found to be green or else not so examined and blue. According to generally agreed inductive common sense, if the emerald inspector examined a long sequence of emeralds before t and found them all to be green, this evidence should make it more plausible than otherwise that an arbitrary emerald examined after t would be found to be green as well. However, this same common-sensical inductive assumption implies that,

if a long series of emeralds were examined before t and found to be grue, this evidence would make it *less* plausible than otherwise that an arbitrary emerald examined after t would be found to be grue. Common sense therefore seems to support an assumption of positive instantial relevance in the case of greenness while supporting an assumption of negative instantial relevance in the case of grueness. In the context of emerald-extraction, then, greenness seems to be projectable and grueness to be non-projectable.

The grue problem illustrates the difficulty of evaluating whether a system of inductive logic is successful without fixing an interpretation of its domain. Suppose that we are interested in evaluating how well the system of inductive logic $(SL_{pred}, \mathcal{M}_{\mathbf{prob},\mathbf{IR}^+})$ formalises the inductive assumptions suggested by common sense in the emerald-extraction situation. Recall that **prob** represents the probability axioms and $\mathbf{IR}^+$ the axiom of positive instantial relevance. Suppose that L_{pred} features only one primitive predicate G and has a large but finite number of constant symbols. If G is interpreted as representing greenness, then $(SL_{pred}, \mathcal{M}_{\mathbf{prob},\mathbf{IR}^+})$ seems to agree with inductive common sense, and all is well. On the other hand, if G is interpreted as representing grueness, then $(SL_{pred}, \mathcal{M}_{\mathbf{prob},\mathbf{IR}^+})$ is at odds with inductive common sense, according to which observing instances of grueness should sometimes make observing further instances *less* plausible. On one interpretation of G, $(SL_{pred}, \mathcal{M}_{\mathbf{prob},\mathbf{IR}^+})$ successfully explicates inductive common sense, whereas on the other interpretation it is unsuccessful.

3.3.3   Artificiality

While one might be tempted to dismiss this problem on the grounds that grueness is a silly and artificial property, this is not a good response. Firstly, as Goodman points out at (Goodman, 1954, p.77-81), grueness is arguably only silly or artificial relative to the ways of describing emeralds that we happen to employ at present. Relative to the emerald-describing method of a Martian for whom grueness is serious and natural, the property of greenness might seem artificial. After all, a given emerald is green if and only if it is either examined before t and found to be grue or else examined after t and found to be 'bleen', that is, observed before t and found to be blue or else not so observed and green. To privilege greenness and blueness over grueness and bleenness without a good reason would perhaps be chauvinistic.

There is another, in my opinion stronger, reason not to dismiss Goodman's objection because grueness is silly and artificial. The reason is that here are many cases of non-projectability that do not involve grueness or any other artificial property. The grue emerald mine is merely one out of many contexts in which common sense supports an assumption of non-positive instantial relevance. Earman (1989) makes this somewhat under-appreciated point as follows:

> ...questions about projectability arise for the most mundane of hy-

potheses and predicates where not the slightest hint of Goodmanian
trickery is present.
(Earman, 1989, p. 220)

Earman's essay does not specify any examples of such mundane contexts: here
are a few. First, consider a person Sam munching their way through a tube of
fruit pastilles.[7] Sam is interested in how likely it is that the last pastille to be
withdrawn from the tube will be blackcurrant flavoured. Nestlé, the company
that manufactures fruit pastilles, say that "we know that everyone has their
favourites which is why we try to make sure all packs include a really good
mix."[8] Sam might take this kind of comment to indicate that it is unlikely
that there will be more than four blackcurrant fruit pastilles in a tube. On this
assumption, Sam might reasonably take the observation of a blackcurrant fruit
pastille at the start of the tube to make future observations of blackcurrants in
the same tube less plausible. In Goodman's terms, Sam considers blackcurrant-
flavour a non-projectable property of fruit pastilles in this context.

Similarly, an ecologist studying tigers might take the observation of a lone tiger
at a particular spot to make future observations of tigers at the same spot less
plausible than in the absence of such an observation, as the observation would
indicate that the spot is likely to lie within the territory of a tiger that is not
raising any children. If the spot had been within the territory of a mother who
was raising some children, then future sightings could be expected to be more
frequent, given the greater number of circulating tigers.

Even the property of greenness is non-projectable in certain contexts. Certain
decks of cards popular in Germany contain a finite number of green cards -
a person repeatedly drawing green cards from such a deck, just like a miner
extracting grue emeralds from the ground, may common-sensically hold that
instances of green before a certain time make instances of green after that time
less plausible.

In sum, the ubiquity of non-projectable properties, together with the fact that
grue is only artificial relative to a certain perspective, show that it is not a
good idea to respond to the grue objection by asserting that non-projectable
properties like grue are silly.

3.3.4   Carnapian inductive logic can formalise assumptions about projectability

Fortunately, Carnapian inductive logic can address the grue problem without
asserting that non-projectable properties like grue should be ignored. The so-
lution is simple: if common sense suggests an assumption of negative instantial

---

[7]Fruit pastilles are cylindrical sweets that come in various flavours. Typically 14 or 15 fruit
pastilles are stacked on top of each other, wrapped in layers of foil and paper and sold as a
'tube'.

[8]This is a quotation from a letter from Nestlé customer services.

relevance with respect to a particular property, then, in order to formalise this assumption within a system of inductive logic, an axiom should be adopted which imposes negative instantial relevance with respect to the predicate that represents that property. Similarly, positive instantial relevance should be imposed with respect to predicates representing projectable properties.

For example, suppose that common sense stipulates only that greenness is projectable in a certain context; say, emerald extraction. One might formalise this assumption as follows. First take a one-predicate unary language L^1_{pred}, interpreted so that its single predicate G represents greenness and symbols $a_1, a_2, \ldots, a_n$ represent successively drawn emeralds. Take as a domain the set of sentences SL^1_{pred}. Next choose a set $\mathcal{M}_{\mathbf{prob,IR^+}}$ of measure functions that is determined by the axioms of probability and positive instantial relevance. The resulting system of inductive logic is therefore $(SL^1_{pred}, \mathcal{M}_{\mathbf{prob,IR^+}})$.

Alternatively, if, in another context, common sense indicates only that greenness is non-projectable, then one could proceed in a similar way. The same domain SL^1_{pred} would be used, and the predicate G would again be interpreted as representing greenness. However, the axiom of positive instantial relevance would not be employed. Since common sense now supports negative instantial relevance with respect to greenness, the axiom of negative instantial relevance must be employed. In this case the resulting system of inductive logic would be $(SL^1_{pred}, \mathcal{M}_{\mathbf{prob,IR^-}})$.

Finally, suppose that, in a third context, common sense indicates that green is projectable and that grue is not. Then a Carnapian inductive logician can proceed by choosing as a domain a new language L^2_{pred} with two predicates: one predicate G representing greenness and a second predicate G' representing grueness. A basic assumption would need to be introduced in order to formalise the logical relationship between greenness and grueness. This could take the form of a stipulation that, for any constant a_i of L_2, the sentence $G'a_i$ is true if and only if either Ga_i is true and $i < t$ or else Ga_i is false and $i \geq t$. The projectability of greenness could be represented by a restricted form of the axiom of positive instantial relevance that applies only to G but not to G'. This would be as follows:

> Restricted axiom of positive instantial relevance
> For the predicate G and any sentence $\theta(a_1, \ldots, a_l) \in SL_{pred}$,
> $m(G(a_{l+2}) \mid G(a_{l+1}) \wedge \theta(a_1, \ldots, a_l)) > m(G(a_{l+2}) \mid \theta(a_1, \ldots, a_l))$

Thus positive instantial relevance would be ensured for the predicate G, which represents a projectable property, but not for the non-projectable G'.

This procedure agrees the spirit of Carnap's last proposal for addressing the grue problem:

> Let us briefly consider a richer language $\mathcal{L}'$ of such a kind that Goodman's abnormal predicates can be defined in it...
> Now it is clear that certain customary inductive procedures that

are valid for 'Blue' and for 'Green' would lead to counterintuitive results if they were applied to the predicate 'Grue' ...

The difficulty may be presented in the following way. Suppose that we are constructing a system of IL for the language $\mathcal{L}'$. We choose one of those procedures or principles which are valid for descriptional attributes like Blue and Green, but invalid for Grue and similar attributes[9]. Suppose we wish to represent this principle by an axiom A of IL. .... The essential point in the formulation of the axiom A is the obvious necessity of including into A a suitable restricting condition, in order to prevent its application to attributes like Grue.

(Carnap, 1980, p. 75-76, footnote added)

Carnap proposes a more sophisticated procedure according to which axioms like the principle of instantial relevance could be restricted on a principled basis. As well ensuring that such axioms do not apply to predicates representing grueness, Carnap's procedure ensures that they would only apply to predicates representing 'descriptional' properties: see (Carnap, 1980, § 4B) for full details.

Thus it seems that Carnapian inductive logic can formalise the assumption suggested by common sense in the case of the grue emeralds.

3.3.5 Why has grue been such a big issue?

Interpretation-neutrality

The widespread belief that the grue problem is a major obstacle for Carnapian inductive logic may be due in part to the common misconception according to which it restricted its attention to interpretation-neutral inductive logical axioms. Such axioms would formalise common-sense inductive assumptions under all interpretations of the domains of the measure functions to which they apply. A system of inductive logic that employed only interpretation-neutral axioms would be 'syntactic' in the sense that it would formalise only those features of common-sense inductive reasoning which depend on nothing but syntax.

The grue problem does indeed show that the axiom of positive instantial relevance is not an interpretation-neutral axiom. If a predicate G is interpreted in an emerald-extraction context as representing grueness, then common sense does not support the assumption that observing the property that G represents makes future observations of the same property more plausible. On the other hand, if G is interpreted as representing greenness, then common sense does support this kind of assumption in the emerald situation. Clearly, this axiom's common-sensicality is interpretation-sensitive.

---

[9]Carnap used the word 'attribute' as we use the word 'property'

However, this fact does not affect the viability of Carnapian inductive logic, because it did not in fact restrict its attention to interpretation-neutral axioms.

This can be seen in general from Carnap's remarks on pure and applied inductive logic. Recall from section 2.4.3 that Carnap thought that the evaluation of systems of inductive logic could only occur within applied inductive logic. By definition, in applied inductive logic, an interpretation is available.

The fact that Carnapian inductive logic did not aim at interpretation-neutrality is also clear from Carnap's final response to Goodman. In this passage he notes that an axiom excluding predicates that represent 'non-descriptional' properties should be evaluated within applied inductive logic, i.e. with respect to an interpretation:

> While this axiom is formulated within pure IL, any argument intended to show its adequacy, and in particular, to show that it excludes nondescriptional attributes like Grue, has to proceed within *applied* IL.
> (Carnap, 1971a, p.76, emphasis original)

Carnap made essentially the same point in his first reply to Goodman, where he argued that inductive logical axioms like positive instantial relevance, which at this stage he called 'inductive rules', should be formulated so as not to apply to properties like grue:

> It seems that those authors who have formulated inductive rules without specifying the kind of properties to which they are applicable (for example, Laplace's rule of succession and Reichenbach's rule of induction) intended them to be applied to purely qualitative properties only [These are essentially the properties that Carnap later referred to as 'descriptional attributes'.]. This was tacitly presupposed; a complete formulation should contain this or a similar restriction.
> (Carnap, 1947b, p.146, round brackets original, square brackets added)

This passage shows that Carnap was happy to distinguish between properties, either 'tacitly' by employing particular axioms in some applications but not others, or explicitly by employing axioms like the restricted principle of instantial relevance that explicitly represent distinctions that common sense may make between different properties.

Richard Jeffrey emphasises that Carnapian inductive logic was not searching for interpretation-neutral axioms in the following passage, written when he was working in close collaboration with Carnap:

> ...Goodman's strictures weighed rather differently upon Hempel's purely syntactical theory than upon Carnap's more semantical one.
>
> Goodman's Query underlined the impossibility of a purely syntactical account of confirmation. No confirmation function can be

attached with equal felicity to all syntactically identical languages. . .

> Then, in choosing a confirmation function to use with a certain language, one must take account of the meanings of predicate expressions that occur in that language. Now I submit that the permanent impact of Goodman's Query on Carnapian confirmation theory consists in having made it abundantly clear that credibilities depend on meanings, so that in choosing a c-function for a language it is essential to consider the meanings of the terms in its vocabulary. To one with Goodman's philosophic conscience, this impact seems deadly, vitiating all further discussion of confirmation in Carnap's terms; but to one who, like Carnap, finds no inherent fault in talk of meaning and necessity, the impact is expected, and salutary.
> (Jeffrey, 1966, p. 282-285)

Despite these clarifications, the misconception arose, and persists, that Carnapian inductive logic was interested only in interpretation-neutral axioms, and that the grue problem was an obstacle for this reason. For example:

> Some Bayesian logicists (e.g., Carnap) have maintained that posterior probabilities of hypotheses should be determined by logical form alone. . . . Most logicians now take the project to have failed because of a fatal flaw with the whole idea that reasonable prior probabilities can be made to depend on logical form alone. Semantic content should matter. Goodmanian grue-predicates provide one way to illustrate the point.
> (Hawthorne, 2014, § 3.2, brackets original)

> Another Goodmanian lesson is that inductive logic must be sensitive to the meanings of predicates, strongly suggesting that a purely syntactic approach such as Carnap's is doomed.
> (Hájek, 2012, § 3.2)

> For Carnap, probability$_1$ is analytic and syntactic. . . . But after his paper proposing a purely syntactic justification for inductive inference [Carnap, 1945b], Nelson Goodman [1946] immediately published a serious challenge to it. Goodman's conclusion was that inductive inference is not purely syntactic in nature. . . Although Goodman and Carnap soon agreed to disagree, there was no escape; and Goodman's point is now generally accepted.
> (Zabell, 2011, p.297-298, square brackets original)

These accounts either claim or imply that Carnapian inductive logic was 'syntactic', by which they seem to mean that it sought to find axioms that were common-sensical under all interpretations of the domains of their measure functions. This misconception partly explains why the grue problem has been seen as a serious obstacle for Carnapian inductive logic.

Not justifying inductive common sense

The solution to the grue problem that I sketch above shows how Carnapian inductive logic can formalise inductive common sense even in 'non-projectable' cases where it does not support assumptions of positive instantial relevance, or even when it supports assumptions of negative instantial relevance. It does not show that inductive common sense is justified. In other words, the old, Humean, problem of induction—that of identifying justified inductive assumptions—remains unsolved.

This is not a serious problem for Carnapian inductive logic. Its primary aim was to formalise inductive assumptions rather than to show that particular assumptions are justified. However, as the discussion of Lakatos's critique showed, it is widely believed that Carnapian inductive logic did in fact aim to find true inductive logical axioms. True axioms, if they existed, would surely formalise justified inductive assumptions. For this reason, someone who share's Lakatos's analysis of its objectives might think that Carnapian inductive logic's response to the grue problem was unsatisfactory.

3.4 Salmon on partial entailment

Salmon (1967) presents another influential objection to Carnapian inductive logic. In this essay Salmon argues that two important inductive assumptions cannot simultaneously be formalised within Carnapian inductive logic.

3.4.1 Hume's principle

The first assumption, which following Jeffrey (1967) I shall call 'Hume's principle', stipulates that 'distinct events' should be 'logically independent'. Salmon explain's Hume's principle as follows:

> [Carnap's] conception faces one severe prima facie problem, namely that, as Hume so forcefully showed, distinct events are logically independent of one another. This is a transparently analytic statement, but it has profound consequences. In particular, it means that the future is logically independent of the past and that the unobserved is logically independent of the observed. How, we must ask, can a probability that is conceived as a logical relation play a fundamental role in inductive logic, if induction is concerned with relations between the observed and the unobserved between the past and the future?
> (Salmon, 1967, p.729)

I shall understand Hume's principle as asserting that propositions that pick out different states of affairs ('distinct events') should neither support nor undermine

one another, whether deductively or inductively ('logical independence'). It is difficult to say whether this is exactly what Salmon meant, as Salmon does not explain at length what it means for events to be distinct or logically independent. However, this reconstruction seems to capture Salmon's drift. Hume's principle can plausibly be formalised within Carnapian inductive logic as 'Hume's axiom':

> Hume's axiom
> If θ and ϕ are logically independent sentences, then $m(\theta \mid \phi) = m(\theta)$

3.4.2   Learning from experience

The second assumption, 'learning from experience', which we have already encountered, stipulates that instances of a certain kind of observation should inductively support further instances of the same kind of observation. Learning from experience can plausibly be formalised by the axiom of positive instantial relevance.

Salmon claims that these two inductive assumptions are incompatible. Fittingly, Hume's axiom is indeed inconsistent with the axiom of positive instantial relevance. To see why, consider the language L^1_{pred} whose only predicate is P, and whose constants are $a_1, a_2, \ldots, a_n$. Suppose that m is a measure function on SL^1_{pred}, and let $\theta = Pa_2$ and $\phi = Pa_1$. Then, if m satisfies the axiom of positive instantial relevance, it should be the case that $m(\theta \mid \phi) > m(\theta)$. However, since θ and ϕ are logically independent, according to Hume's principle it should also be the case that $m(\theta \mid \phi) = m(\theta)$. Clearly the two axioms are incompatible.

3.4.3   Salmon's argument

According to Salmon, Hume's principle is obviously a true inductive assumption, whereas learning from experience is not. He therefore concludes that the incompatibility of the two assumptions shows that learning from experience must be a false inductive assumption. Consequently, systems of inductive logic that formalise only true inductive assumptions, or in Salmon's terms 'capture partial entailment', should rule out all measure functions that satisfy positive instantial relevance:

> Unfortunately for induction, statements strictly about the future (unobserved) are completely independent of statements strictly about the past (observed). Not only are they deductively independent of each other, but also they fail to exhibit any partial entailment. The force of Hume's insight that the future is logically independent of the past is very great indeed. It rules out both full entailment and partial entailment. If partial entailment were the fundamental concept of inductive logic, then it would in fact be impossible to learn from experience. If instead, like Carnap, we adopt

another confirmation function (such as $c^\star$) that permits learning from experience, then we have abandoned the concept of partial entailment as the basis for inductive logic.
(Salmon, 1967, p.731-732)

3.4.4   Responses to Salmon's objection

The most important repsonses to Salmon's critique are Jeffrey (1967) and Kemeny (1967).

Jeffrey's response questioned the intelligibility of Hume's principle. According to Jeffrey, Salmon did not specify Hume's principle sufficiently clearly for it to have important consequences for inductive logic:

Now Hume's principle, and Salmon's, will be intelligible or not, and will (if intelligible) be true or false, depending on how the following terms are understood: event, logical independence of events, and reference of sentences to events. None of these terms play any role in logic, where logical independence is a relation between sentences. I know of no standard usage for them and I conclude from the fact that Professor Salmon uses them without explanation that passages like the two I have quoted [Jeffrey quotes the passage from (Salmon, 1967, p.729) and another passage discussing Hume's principle from later in Salmon's article.] are properly viewed not as parts of an argument against the logical interpretation of probability, but rather as intimations of the sort of thing that troubles him.
(Jeffrey, 1967, § 0, square brackets added)

While Jeffrey was right to point out that Salmon had not expressed Hume's principle as clearly as he might have, I think that Jeffrey's response is too quick. An inductive logical axiom, namely Hume's axiom, *can* be extracted from what Salmon says about Hume's principle by equating 'distinct events' with logically independent sentences and the inductive independence of distinct events with the condition that Hume's axiom imposes. The incompatibility of Hume's axiom and positive instantial relevance bears out Salmon's fundamental worry.

Kemeny (1967) shows how a better response to Salmon's objection might proceed. While the main body of Kemeny's essay focusses on the unrelated question of how to formalise reasoning about scientific theories within Carnapian inductive logic, Kemeny briefly touches on the issue of Hume's principle in this passage:

First of all, I should like to point out that Salmon discusses three (not two) different problems: (1) Formulation of an inductive logic, which we will identify with the choice of a confirmation function. (2) The selection of a rule of induction.[Kemeny uses the term 'rule' as

we use the term 'axiom'] (3) The justification of induction in general, or the justification of a particular inductive rule.

I am convinced that Hume has spoken the definitive words concerning the third problem. One cannot justify induction. One can at best proclaim it as an act of faith or give pragmatic reasons for accepting it as a guide in life. I would, therefore, prefer to return to Carnap's original approach. Let us look not at the problem of justification, but at the problem of explicating the method by which scientists predict the future. The task of the philosopher is to make clear what the scientist does, and to try to improve on it by substituting precise tools for a somewhat nebulous intuition.
(Kemeny, 1967, p. 2)

As Kemeny points out, although it is difficult, if not impossible, to justify the assumption that observed instances make unobserved instances more plausible, it can still be useful to formalise this inductive assumption. Carnapian inductive logic is concerned mainly with whether inductive assumptions can usefully be formalised, and not with whether or not they are true. Thus, whether or not they agree with Salmon that Hume's principle is 'transparently analytic', Carnapian inductive logicians are free to define systems of inductive logic that do not validate it.

It might seem like Salmon's criticism is very different from that of Lakatos. This is true as far as their conclusions went: whereas Lakatos concluded that Carnap's whole programme was degenerate, Salmon drew only the more measured conclusion that Carnapian inductive logic should not employ axioms that are inconsistent with Hume's axiom. However, the reasoning that led Salmon and Lakatos to their different conclusions was surprisingly similar. Both authors' arguments assume that Carnapian inductive logic aimed to find true, theoretically justified inductive logical axioms, whereas in fact it only aimed at pragmatic justification. When we discard this mistaken assumption, we see that Carnapian inductive logic survives both critiques.

Curiously, although he makes the same error as Lakatos, Salmon's historical account is the reverse. Rather than claiming, like Lakatos, that Carnapian inductive logic began by searching for true axioms and then settled for mere usefulness, Salmon claims that Carnapian inductive logic began by formalising and ended by aiming for truth:

If Carnap were content, as he was twenty years ago, to claim that his sole purpose is to explicate our intuitive sense of what constitutes correct induction, without trying to provide any actual justification for it, then one could hardly have any serious misgivings about the approach via intuition. If, on the other hand, conformity with intuition is taken as providing justification as well, as Carnap has more recently maintained, fundamental problems do arise.
(Salmon, 1967, p. 737-738)

It is unclear why Salmon took Carnap to have maintained that inductive intuition can provide justification: fortunately we do not need to do so.

3.5 Other critiques of Carnapian inductive logic

3.5.1 Universal hypotheses

According to an objection frequently articulated by Popper, for example in (Popper, 1959, § 80), Carnapian inductive logic was unsuccessful because it could not capture reasoning about universal hypotheses such as 'all ravens are black'.

Unary predicate languages can represent universal hypotheses using universally quantified sentences: for example, the sentence $\forall x(Rx \to Bx)$ might represent the universal law according to which all ravens are black. However, the measure functions allowed by Carnap's favoured systems of inductive logic assign all such sentences the value 0, provided that the relevant language has countably infinitely many constant symbols, as seems appropriate when formalising hypothesis that does not refer to a specific number of ravens. Consequently, it is difficult to represent reasoning about universal hypotheses using this kind of system of inductive logic.

Carnap's response to this kind of objection was to argue that, on the one hand, it would be interesting to investigate systems of inductive logic that did not have the property of forcing their measure functions to assign the value 0 to universally quantified sentences of languages with infinitely many constant symbols. Carnap makes a comment to this effect at (Carnap, 1963b, p. 976) This ambition was realised by the work of Hintikka, who developed systems of inductive logic in which measure functions assigning positive values to such sentences are not excluded. This work is summarised in Niiniluoto (2011).

On the other hand, Carnap argued, his systems of inductive logic did not need to be capable of representing reasoning about universal hypotheses in order to be useful. While his systems could not describe reasoning about hypotheses themselves, they could describe reasoning about specific instances of universal hypotheses (Carnap, 1950b, § 110G), and hypotheses themselves can often be done without (Carnap, 1950b, § 110H). In addition, Carnap noted, just as deductive logic is useful, despite the fact that not all deductive reasoning that scientists carry out is formalised, so inductive logic need not formalise all aspects of scientific reasoning. He put the point as follows:

> The situation in inductive logic is similar [to the situation in deductive logic]. ... there are many situations which, by their complexity, make the application of inductive logic practically impossible. For instance, we cannot expect to apply inductive logic to Einstein's general theory of relativity.... The same holds for the

other steps in the revolutionary transformation of modern physics,
especially in connection with the quantum theory. ...

On the other hand, there are also cases in which there are good
reasons for the expectation that the application of inductive logic will
become useful for the scientist, or in which the useful application is
possible today.
(Carnap, 1950b, p. 243)

I think that this last response is by far the strongest to the objection based
on universal hypotheses. While statements like 'all ravens are black' may not
seem complex in comparison to the general theory of relativity, the same point
applies. Even systems of inductive logic which cannot describe reasoning about
any universal hypotheses may still usefully formalise other interesting aspects
of scientific reasoning.

3.5.2   Constant exchangeability

Many critiques of Carnapian inductive logic have taken the form of objections
to the axiom of constant exchangeability. Various authors have argued that
constant exchangeability, unlike weaker inductive logical axioms, is questionably
logical. The nature of this kind of objection can be discerned from the following
quotations:

At this point, it is important to ask, In what sense are Carnap's
theories of logical probability (especially his later ones) *logical* ? His
early theories... applied something like the principle of indifference
to the state and/or structure descriptions of the formal language
L in order to determine the logical probabilities $P(\bullet \mid \bullet)$. In this
sense, these early theories assume that certain sentences of L are
equiprobable a priori. Why is such an assumption *logical*? Or, more
to the point, how is *logic* supposed to tell one which statements are
equiprobable a priori?
(Fitelson, 2005, p. 390, emphasis original)

'logical' probability measures, whether based on the Principle of
Indifference or on some other method of distributing probabilities
a priori, do not, we believe, possess a genuinely logical status. For
such systems are ultimately quite arbitrary, and we take logic to be
essentially non-committal on substantive matters.
(Howson and Urbach, 1989, p. 72-73)

Significantly, Carnap's various axioms of symmetry are hardly
logical truths.
Hájek (2012)

To some extent these objections merely reiterate Lakatos and Salmon's view
that inductive logical axioms should be true—note that Hájek uses the term

'truth' explicitly whereas the other two authors use the truth-related term 'a priori'—and claim that constant exchangeability is not true, or at least not true a priori. This line of criticism can be overcome by pointing out, again, that Carnapian inductive logic did not aim to find true axioms.

On another reading, these criticisms are not so concerned about truth, but rather assert that, whether or not it is true, constant exchangeability goes too far to be genuinely logical. On this reading, the objections can be overcome by noting that, while some conceptions of logic may include weaker axioms but exclude constant exchangeability, Carnap's conception certainly did not. According to Carnap's tolerant conception, as we saw in section 2.9.2, *any* axiom for measure functions is logical in the sense that it can be part of a system of semantic rules specifying what follows from what. Even if Carnap's conception of logic were ultimately shown to be incorrect, his research programme could still proceed exactly as before, though without the label 'logic'. Rather than dwelling on the issue of whether particular axioms are or are not genuinely logical, it seems preferable to address the question of whether constant exchangeability and other axioms play the role that they are supposed to within Carnapian inductive logic: that of usefully formalising inductive assumptions.

Another line of criticism of constant exchangeability does address its usefulness, asserting that the assumption that constant exchangeability formalises—loosely that the order of experiments should be inductively irrelevant—is sometimes inappropriate and therefore not useful to formalise. Specifically, this is the case whenever inductive logic is being used to model reasoning in situations where it is appropriate to take the order of experiments into account.

Such situations happen fairly often: to take one, consider Williamson's example from Williamson (2015) of a series of experiments recording the results of the 'game of red and blue':

> In this game, a fair coin is tossed, changing a score s, which is initially zero, to $s + 1$ if heads occurs or $s - 1$ if tails occurs. If $s \geq 0$ the result of the toss is blue, if $s < 0$ the result is red.
> (Williamson, 2015, Ch. 4)

It would almost always be appropriate to take into account the order of experiments in this case, as the order of results of this game can be highly informative. For example, if one 'blue' result is followed by two 'red's, this implies that $s = -2$, meaning that it is impossible for the next result to be 'blue'. The same results in a different order—say 'red', 'blue', 'red'—would not have the same implication. As a result, it would not be appropriate to impose constant exchangeability when formalising assumptions about the game of red and blue.

Carnap was well aware of this issue, as can be seen from the following quotation where he acknowledges that constant exchangeability, which he refers to as the 'principle of symmetry', is not always appropriate:

> In some series of events, the temporal order has no influence. For

example, we find that, in a long series of throws of any die, even if it is loaded, the relative frequency of an ace among those throws that follow immediately on an ace (or a deuce), is not essentially different from the relative frequency of an ace in the whole series. Thus experience shows here statistical independence; and therefore we treat the members of such a series as inductively independent. In contrast, we find a strong dependence in a series of meteorological observations made at a fixed place at noon of every day; and still higher dependence if made at intervals of one hour. The relative frequency of rain immediately following on rain is much higher than in the whole.

If now in elementary inductive logic, we decide to accept the principle of symmetry for a certain investigation, this decision may either be based on the assumption that statistical independence actually holds; or else on the assumption that there is some dependence, but so small that we may neglect it for the sake of simplicity; or finally on the assumption that there may be noticeable dependencies in the actual temporal order represented in another (unknown) language L', and that the order of indexes in our language L was produced by an (unknown) permutation π of the indices in L', where π is an infinite permutation with bounded displacement.
(Carnap, 1971a, p.120)

Although constant exchangeability is not appropriate for all applications of Carnapian inductive logic, this fact does not undermine the programme. Carnapian inductive logic did not make any hard-core commitments involving constant exchangeability, and certainly was not committed to imposing this axiom in all circumstances. Constant exchangeability featured only in secondary claims, according to which certain axioms usefully explicate certain inductive assumptions.

In order for such claims to be warranted, it would only need to be the case that constant exchangeability is a useful axiom in some circumstances: this possibility is not precluded by the fact that it is sometimes not useful.

Despite this fact, given its technical importance and prominence within Carnapian inductive logic's protective belt, we might justifiably demand that constant exchangeability be a useful axiom for *many* applications of inductive logic.

This is surely the case: there are many situations where formalising inductive assumptions is called for and where it is appropriate to assume that the order in which experiments are carried out does not convey any information. As Carnap argued in the quotation above, such situations include those where statistical independence is judged to hold, whether fully or approximately, as well as those where it is not clear how the order of experiments is significant. This kind of situation is very common in statistical research. Gelman et al. go so far as to claim that it is typical for statistical investigations to begin with an assumption that amounts to constant exchangeability:

> The usual starting point of a statistical analysis is the (often tacit) assumption that the n values y_i [representing experiments] may be regarded as *exchangeable*.
> (Gelman et al., 1995, p.5, emphasis original, square brackets added)

Given how often the kind of assumption that it formalises is made in practice, it is clear that constant exchangeability is a useful inductive logical axiom. Critiques of Carnapian inductive logic centred around objections to constant exchangeability are therefore unsuccessful.

3.5.3 Reliability in the limit

Putnam (1963) contains an argument against Carnapian inductive logic that is structurally similar to Salmon's. Like Salmon, Putnam presents a putatively essential inductive assumption, namely 'reliability in the limit', and argues that Carnapian inductive logic cannot easily formalise it. Putnam's criticism is more ambitious than Salmon's, however. Whereas Salmon claims that formalising Hume's principle frustrates Carnap's attempt to capture learning from experience, Putnam aims to show that the difficulty of formalising reliability in the limit shows that the whole enterprise of representing plausibility using measure functions is ill-founded. In addition, Putnam claims that an alternative approach to formalising inductive reasoning does not suffer from this problem.

Reliability in the limit is an inductive assumption which asserts that, based on evidence according to which sufficiently many observations agree with a given hypothesis about the whole sequence of observations, it should be more likely than not that the next observation in the sequence will also agree with the hypothesis. In other words, after enough observations, it should be possible to 'learn' any pattern in a sequence at least to the extent that all subsequent manifestations of the pattern are plausible to a degree strictly greater than $1/2$.

Putnam claims at (Putnam, 1963, p.766) that failing to accept reliability in the limit as a reasonable inductive assumption amounts to the assertion that there are certain true hypotheses that science is unable to accept permanently. Such a claim, Putnam says, would be false, given the availability of the method of 'trial and error'. According to this method, hypotheses are sequentially chosen arbitrarily, accepted and then tested against incoming data until either refuted or accepted forever. This method, Putnam argued, eventually accepts any true hypothesis permanently.

In order to formalise reliability in the limit within Carnapian inductive logic, Putnam introduces some new notions, which I summarise here using the terminology of this thesis.

A *data sequence* $e_{k,\vec{i}}$ with length $k \in \mathbb{N}$ and characteristic vector $\vec{i} \in \mathbb{N}^k$ is a sentence $\alpha_{i_1}(a_1) \wedge \ldots \wedge \alpha_{i_k}(a_k)$ asserting that constants $a_1, \ldots, a_k$ instantiate atoms $\alpha_{i_1}, \ldots, \alpha_{i_k}$.

An *effective hypothesis* is a countably infinitely long data sequence.

Based on these definitions reliability in the limit can be defined as follows, as an axiom applying to measure functions $m : SL_{pred} \to \mathbb{R}$, where L_{pred} is a unary predicate language with countably infinitely many constants:

> Axiom of reliability in the limit: **RL**
> For every possible effective hypothesis h, there is a number $r \in \mathbb{N}$ such that, for all data sequences $e_{k,\vec{i}}$ such that $k > r$, if $\vDash h \to e_{k,\vec{i}}$ and $\vDash h \to \theta(a_{k+1})$, then $m(\theta(a_{k+1}) \mid e_{k,\vec{i}}) > 1/2$.

There are no measure functions of the right kind satisfying the axiom of reliability in the limit. This is shown at (Putnam, 1963, p.768-769) and (Kelly et al., 1994, p.6-7). This seems to show that Carnapian inductive logic cannot formalise an important inductive assumption.

To make matters even worse for Carnapian inductive logic, Putnam proposed an alternative formal method which, he argued, could capture reliability in the limit. This method, which Putnam called M, attempts to formalise the method of trial and error, according to which hypotheses are exhaustively accepted and then rejected them if and when they disagree with incoming data. Putnam compares M with acceptance procedures based on Carnapian inductive logic, noting that someone using M "can do things... that he could not *in principle* accomplish if he [used] degree of confirmation"(Putnam, 1963, p.773, emphasis original).

If Putnam is correct that the axiom of reliability in the limit should be satisfied if possible, then this seems to show that Carnapian inductive logic should be rejected in favour of his method M in such cases. I argue below that reliability in the limit is not, in fact, a useful axiom.

Carnap's response

Carnap (1963c) disputes Putnam's argument for three separate reasons. First he claims that reliability in the limit cannot be satisfied by measure functions that take into account other factors than compatibility with the evidence, such as, for example, simplicity. Such measure functions never guarantee permanent acceptance of true hypotheses, Carnap claims, because at any point a true but bad-with-respect-to-the-extra-factor hypothesis might be rejected in favour of a false hypothesis that is also compatible with the available evidence. Since scientific reasoning takes into account such extra factors, Carnap concludes, it would be unreasonable to insist that measure functions that formalise scientific reasoning satisfy the axiom of reliability in the limit.

This objection strikes me as plausible, but it is hard to judge definitively because simplicity and other theoretical virtues are difficult notions to pin down.

Carnap's second objection to reliability in the limit contests the assumption,

implicit in Putnam's defence, that systems of inductive logic should be judged according to their long-run properties. Such properties, he argues at (Carnap, 1963c, p.985), have little correlation with practical usefulness.

I think that this objection is potentially strong: a system of inductive logic might be very useful for everyday purposes, where only data strings of length up to say, 100 trillion, need to be considered, and yet begin to be problematic at the 100 trillionth iteration. On the other hand a system with very good long-run properties might be useless in every way until it reaches the 100 trillionth iteration. Axioms such as reliability in the limit therefore seem only loosely related to its usefulness. However, Carnap was prepared to consider other axioms that refer to long-run properties, such as Reichenbach's axiom, which he discusses at (Carnap, 1980, § 20). As a result some additional argumentation would be required to show a difference between good and bad long-run axioms. I will not consider this question as there are alternative counterarguments to Putnam.

Carnap's final objection to Putnam asserted that inductive logic should not seek to guide decisions whether or not to accept hypotheses in the way that the axiom of reliability in the limit assumes. Implicit in the axiom is the idea that there should be some threshold t greater than $1/2$ such that, if the m-value of a sentence is greater than t, then m should be thought of as 'accepting' the sentence. Carnap put his objection to this kind of 'rule of acceptance' as follows:

> In my view, rules of this kind give in some respect too much, in another respect too little. I shall briefly indicate my main reasons for this opinion.
>
> (a) Suppose that rules of acceptance are given which on the basis of the total evidence e available to X determine the acceptance of some hypothesis h. This means practically that the rules tell X to act as if he knew that h were true. But such an action may be entirely unreasonable. Therefore the rules say more than they should say. It is impossible to give rational advice for practical action merely on the basis of logical relations between e and h; for this purpose the expected gains or losses (more exactly, their utilities for X) must also be taken into account.
>
> (b) In contrast, there are certain situations in which the rules of acceptance would not provide X with any advice on how to act. Suppose, for example, that X knows that 100 balls have been drawn from an urn containing black and white balls, and that among them were 60 black balls and 40 white balls. X is allowed to choose one of the acts $a1$ and $a2$; if he chooses $a1$ he obtains \$100 if the next ball is black ($H1$), and otherwise nothing; similarly for $a2$ with white ($H2$). Since on the basis of e neither $H1$ nor $H2$ can be predicted with reasonable confidence, rules of acceptance of the customary kind will leave both hypotheses in suspension and therefore will not recommend either of the two acts. But obviously it would be rational for X to choose $a1$.

(Carnap, 1963c, p. 971-972)

Thiss objection to the axiom of reliability in the limit seems reasonable. It is not clear that measure functions should guide acceptance rules in the way that Putnam supposes.

Another response

I think there is another telling objection to Putnam's critique that Carnap did not identify. Putnam failed properly to distinguish the method M, which he claimed is good, from measure function based methods, which he claimed are bad.

Kelly, Juhl and Glymour show at (Kelly et al., 1994, p.7-16) that reliability in the limit not only rules out acceptance methods based on Carnapian measure functions, but in fact rules out all computable hypothesis-acceptance methods. Putnam's method M, they show, manages to formalise reliability in the limit precisely because it is non-computable, requiring an 'oracle' to ensure that the series of hypotheses it receives is both complete and free of infinite loops. Furthermore, they show (Kelly et al., 1994, p.11) that, if Carnapian inductive logic had been afforded a similar privilege, it would have been possible to define measure functions that satisfy the axiom of reliability in the limit.

The comparison that Putnam invites is therefore unfair: the fact that M succeeds where Carnapian inductive logic fails owes more to the oracular advantage that Putnam affords his method than to any failing of Carnapian inductive logic.

Conclusion

In sum, Putnam's argument against Carnapian inductive logic was unsuccessful. In light of Carnap's second objection it is not clear that it is useful to formalise reliability in the limit. Even if formalising reliability in the limit were a useful goal, Kelly, Juhl and Glymour's demonstrations show, contrary to Putnam, that Carnapian inductive logic is as good a tool for the job as Putnam's M.

3.6   Discussion

In this chapter I considered critiques of Carnapian inductive logic by Lakatos, Goodman and Salmon, as well as several more minor critiques. I argued that none of these critiques succeeded in showing that Carnap's programme failed in any significant way.

I have not touched on every prominent critique of Carnapian inductive logic. In particular I have not discussed a dispute between Carnap and Popper over

whether or not measure functions should be interpreted as representing 'confirmation'. According to Popper, an objection arising out of this question shows "that Carnap's theory is self-contradictory, and that its contradictoriness is not a minor matter which can be easily repaired, but is due to mistakes in its logical foundations." (Popper, 1959, p. 409) I have not been able to reconstruct Popper's argument so as to make it amount to anything more than a terminological objection. Readers who hope to get to the bottom of the dispute are referred to (Popper, 1959, Appendix 9), Popper (1968) and Carnap's replies to Popper in Lakatos (1968c) and Carnap (1963e).

3.6.1 Why were there so many critiques?

It is interesting and puzzling that critiques of Carnapian inductive logic have been so abundant and diverse. Why did so many people think that Carnap was wrong for so many different reasons?

Carnapian inductive logic might have provided such a hospitable environment for critiques because of poor communication on Carnap's part. All of the critiques depend crucially on misconceptions about Carnapian inductive logic: for example that it aimed at truth, proposition-neutrality, long-run reliability or an explanation of scientific theory-change, or that it sought to show that inductive assumptions of independence or learning from experience are theoretically justified.

These misattributions are understandable because Carnap's writing on inductive logic can be very difficult to interpret. Carnap used familiar terms like 'objective', 'logical', 'analytic' etc in un-orthodox and suggestive ways that encourage misreading. Often—for example in the cases of Carnap's remarks on the blank slate situation, the objectivity of probability$_1$ and pure and applied inductive logic—a very careful reading is required to work out exactly what Carnap meant. It seems that these issues combined to make it seem to secondary authors that Carnap had made some obvious mistakes.

The importance of Carnap's misleading writing is emphasised by the fact that some of the only secondary accounts of Carnapian inductive logic that did not make interpretative errors were written by Jeffrey and Kemeny. See, for example Jeffrey (1975), Jeffrey (1966), Jeffrey (1973), Jeffrey (1967), Kemeny (1967), Kemeny (1955), Kemeny (1963b) and Kemeny (1963a). Both of these authors worked closely with Carnap in person over many years, giving them, unlike other commentators, an alternative way to understand his views.

The task of interpreting these passages has been made much easier by some relatively recent developments in the secondary literature on Carnap. Carnap's systems of inductive logic can now be understood at a mathematical level far more easily than has been the case in the past thanks to the work of the Manchester school, which has presented the formal aspects of Carnapian inductive logic

in a modern and succinct way. In contrast, Carnap's own mathematical presentation is idiosyncratic, with unusual symbols that are difficult to decipher for the modern reader, and interspersed with philosophical commentary, increasing the effort required to read either.

In addition, Carnap's overall approach to philosophy has been greatly clarified by the many recent efforts towards a reappraisal of his overall approach to philosophy, which I mentioned at the start of chapter 1. A proper understanding of Carnap's overall approach was therefore not available to many of the authors who criticised Carnapian inductive logic.

3.6.2   What next?

Let us now assess where we stand. This thesis aims to show that Carnapian inductive logic merits reappraisal and begin the process of reappraising it, both historically and substantively. The previous chapter construed Carnapian inductive logic as a research programme, argued that it had many of the features that Lakatos saw as indicating success and addressed some interpretative points. This chapter claimed that the most influential critiques of Carnapian inductive logic were unsuccessful.

A reader who agrees with these arguments will recognise that Carnapian inductive logic has at least two attractive features: it succeeds according to Lakatos's criteria and survives its most prominent critiques. However, the reader might wonder whether there are might be other rival research programmes that tackle the same problems as Carnapian inductive logic more successfully, and whether Carnapian inductive logic has any practical applications.

The next chapter begins to address the first of these queries, arguing that one rival research programme, namely subjective Bayesian inductive logic, is less successful in some respects than Carnapian inductive logic.

Chapter 4

Subjective Bayesian inductive logic

4.1 Introduction

If there were a rival research programme which is so much more successful than Carnapian inductive logic as to render it redundant, then the case for substantive reappraisal would be weakened. Carnapian inductive logic might be interesting from a historical point of view, but would not have much to offer philosophers interested in reviving it. On the other hand, if Carnapian inductive logic had any pronounced advantages over its rivals, this would strengthen the case. In order to find out whether to substantively reappraise Carnapian inductive logic, we must therefore investigate how it compares with rival research programmes.

The next two chapters compare Carnapian inductive logic with two rival research programmes—subjective and objective Bayesian inductive logic—arguing that Carnapian inductive logic has advantages over both, thanks to its comparatively uncontentious philosophical presuppositions and methodological flexibility. Whereas both objective and subjective Bayesian inductive logic are committed to bold and somewhat dubious epistemological positions, Carnapian inductive logic leaves most epistemological issues open. Similarly, both objective and subjective Bayesian inductive logic stipulate that evidence should be represented in certain ways but not others, whereas Carnapian inductive logic makes neither of these methodological prescriptions.

Subjective Bayesian inductive logic proposes to formalise inductive assumptions in a way that is consistent with the tenets of subjective Bayesian epistemology. In this chapter I introduce subjective Bayesian epistemology and argue that two of its epistemological commitments are dubious, whereas Carnapian inductive logic does not make any similarly contentious epistemological claims. I take

these conclusions to demonstrate that Carnapian inductive logic is preferable to subjective Bayesian inductive logic, at least as far as its philosophical presuppositions are concerned.

I also consider which form of inductive logic is preferable from a methodological point of view. Here the picture is more complicated. On the one hand, subjective Bayesian inductive logic has the advantage of being comparatively simple. On the other hand, Carnapian inductive logic has a wider range of applications and is not limited to representing evidence using only the conditional probability method. I argue that, on balance, these methodological considerations favour Carnapian inductive logic.

I conclude that, due to its philosophical and methological advantages, Carnapian inductive logic is preferable to subjective Bayesian inductive logic.

4.1.1   Inductive logic and epistemology

This chapter and the next discuss various connections between inductive logical and epistemological research programmes. It will be useful to discuss the relationship between inductive logic and epistemology in general before looking at specific cases.

Inductive logic and epistemology differ principally in the type of claim that they aim to produce. Whereas epistemological research programmes aim to produce claims about the nature of rational belief, inductive logical research programmes aim to produce formal representations of inductive reasoning. The two disciplines are very closely related, as work in the epistemology of partial belief often relies on idealised formal representations of inductive reasoning, while inductive logic often focuses on formalising the inductive assumptions of rational agents. Nonetheless, it is possible in principle to do one without doing the other: inductive logic need not concern questions of rationality, and rational partial belief can be studied without the use of formal representations.

I do not wish to defend the view that any discipline that aims to formalise inductive reasoning is an inductive logic. There are many different conceptions of inductive logic, not all of which are so liberal. We have seen that Carnap's conception is extremely liberal, allowing that any system of semantic rules according to which some claims about inductive reasoning follow from others can, in principle, amount to an inductive logic. Colin Howson's conception is more restrictive, stipulating that genuine inductive logics must produce consistency constraints for degrees of partial belief. According to Williamson's conception, as presented in Williamson (2013), inductive logic must specify entailment relationships with the form $\phi_1^{X_1}, \ldots, \phi_k^{X_k} \approgt \psi^Y$, where $\phi_1, \ldots, \phi_k$ and ψ represent propositions and $X_1, \ldots, X_k$ and Y represent inductive qualities.

Despite the differences between these conceptions, in each case inductive logic aims to produce formal representations of inductive reasoning, and the task of

doing so is conceptually distinct from the epistemological task of describing the nature of rationality. Carnap was clear that formalisation was the principal aim of his research programme, and explains in Carnap (1971b) his view about the relationship between inductive logic and claims about rationality, according to which the two are distinct. Howson emphasises the same distinction in the following passage:

> A set of constraints determining consistency does not by itself assume that anyone does or even can satisfy them in every instance. To weaken them so that they are all humanly achievable is a misguided enterprise, rather like weakening the deductive principle that a contradiction implies everything to one stating merely that known contradictions imply everything.
> (Howson, 2000, p.151)

Howson thinks that the epistemological phenomenon of rationality must be humanly achievable, whereas the inductive logical phenomenon of consistency need not: the two must therefore be distinct.

Williamson's conception stipulates that inductive logic must specify entailment relationships that represent valid inductive inferences. Again, according to this conception, inductive logic is in the business of producing formal representations of inductive reasoning and not necessarily that of identifying which states of partial belief are rational.

By way of comparison, both deductive logicians and epistemologists who study the rationality of states of full belief are careful to distinguish these two enterprises. The task of investigating the nature of rational full belief is typically considered part of epistemology, whereas the task of constructing formal representations of deductive reasoning is considered part of logic. Just as in the inductive case, these two tasks are closely, but not necessarily, related, and remain distinct according to various conceptions of deductive logic.

4.1.2 Subjective Bayesian epistemology

Subjective Bayesian inductive logic stems from subjective Bayesian epistemology, a prominent philosophical programme which makes claims about the conditions under which an agent's state of partial belief is rational.

The most important claims that subjective Bayesian epistemology makes are called probabilistic necessity, probabilistic sufficiency and conditionalisation. Probabilistic necessity and sufficiency assert that, in order for a state of partial belief to be rational, it is respectively necessary and sufficient that it can be represented by a probability function. Conditionalisation asserts that, in response to obtaining novel evidence, rational agents update their states of partial belief using a certain method. Specifically, a rational agent whose pre-evidence state of belief is represented by the probability function $pr_{initial}(\bullet)$ should ensure that

the probability function $pr_{new}(\bullet)$ representing their post-evidence state of belief is the conditional probability function $pr_{initial}(\bullet \mid E)$, defined in the standard way as set out in section 2.4.2 above, where E is a member of the domain of $pr_{initial}$ that represents the evidence.

Perhaps the most famous subjective Bayesian was Bruno De Finetti. His position with respect to probabilistic necessity and sufficiency can be seen from the following quotation:

> Our goal is, in essence, to characterise the set of formally admissible opinions. We are not concerned with the possibility of there being another criterion which would allow us to consider one opinion to be more or less correct than another. Those criteria would in fact go beyond the purely logical aspect of the problem which can, and should, be tackled only by mathematics.
>
> As a consequence, a strict separation of two moments appears to be necessary: the characterisation of non-incoherent opinions, the formal moment which must be dealt with mathematically; the choice of one among those possible opinions which must be left to practice, common sense and the judgement of the individual.
>
> The only difference between those who subscribe to the subjective point of view and those who subscribe to the objective one is the following: whereas subjectivists consider such a practical choice to be free and arbitrary, objectivists think that it can be made correctly in exactly one way.
> (De Finetti, 1931, p.299)[1]

De Finetti thought that non-probabilistic states of partial belief, which he calls 'opinions', are 'formally inadmissible': they are ruled out by formal rationality conditions. Choices between states of partial belief that are not so ruled out, on the other hand, cannot be guided by principles but only by "free and arbitrary" choices guided by informal common sense and judgement. In other words, representability by a probability function is necessary and sufficient for a state of partial belief to be rational.

Subjective Bayesian epistemology continues to be debated in the formal epistemological literature, and has been championed in various different forms. See Joyce (2011) for an overview.

Howson's subjective Bayesianism

Colin Howson is a prominent contemporary commentator on subjective Baysesian epistemology, and, I shall now claim, is best thought of as a defender of it.

---

[1]Translation very kindly provided by Hykel Hosni, to whom I am also grateful for bringing the passage to my attention.

Howson sometimes writes that, whereas the probability axioms are, in his view, "laws of consistency", they do not describe "rationality". For example, at (Howson, 2003, p. 154), he claims that it is "misguided" and a "misidentification" to think of the probability axioms as providing a "model of rational belief", and that the "standard Bayesian model" does not provide a "plausible theory of rationality". In Howson (2000), he takes a more guarded but similar position, giving the impression that rationality is not his main focus:

> I have scrupulously avoided discussing scientific rationality, partly because it is a highly contested area, but mainly because this is a book about logic, not about rationality.
> (Howson, 2000, p.239)

It might be thought, given these remarks about rationality, that Howson is not committed to subjective Bayesian epistemology. Nonetheless, I believe that Howson is best viewed as a subjective Bayesian epistemologist, albeit a somewhat atypical one. This is because Howson appeals to another concept— 'fairness'—which is functionally very similar to the epistemological concept of 'rationality', at least as far as it appears in the literature on subjective Bayesian epistemology.

The main claim of Howson (2003) is that the probability axioms are fairness constraints for 'betting quotients', which he regards as suitable measures of degrees of partial belief:

> How should one measure degrees of belief? The Bayesian literature contains a number of suggestions. It would take too long to review all of them here. What I shall do is take one, and show how the probability axioms emerge as the appropriate constraints. The measure I will consider is the individuals *fair odds*, or to be more precise, their *fair betting quotient*.
> (Howson, 2003, p. 159, emphasis added)

Howson therefore sees the probability axioms as connected with states of partial belief, via betting quotients. Moreover, it is a consequence of his approach that representability by a probability function is necessary and sufficient for a state of partial belief to be measured by a fair betting quotient. Just as subjective Bayesian epistemologists see representability by a probability function as necessary and sufficient for a state of partial belief to be rational, Howson sees it as necessary and sufficient for a measure of partial belief to be fair. Howson also endorses a fairness-based form of conditionalisation: see (Howson, 2003, § 5). Howson is even comfortable with evaluating reasoning according to whether or not it is probabilistic. In Howson (2000), he claims many times that the probability axioms provide standards of "sound" or "correct" reasoning. Fairness therefore plays a very similar role within Howson's set-up to that of rationality within subjective Bayesian epistemology.

Howson's objection to the term 'rational' seems not to stem from a conviction that the probability axioms are inappropriate standards for states of partial

belief, but rather from a specific view about rationality, compared to other standards. Howson clearly believes that rationality, unlike measurability by a fair betting quotient, should be achievable by real agents. Since, in his view, representability by a probability function is unachievable, Howson concludes that the probability axioms provide standards of fairness, but not of rationality.

In order to do full justice to Howson's views, it would perhaps be best to categorise him as a fairness theorist rather than a subjective Bayesian epistemologist, while addressing in detail the questions of whether rationality and fairness really should be achievable by human agents. However, given the similarity between subjective Bayesian standards for evaluating states of partial belief and Howson's, it seems best for our purposes to classify him as a subjective Bayesian epistemologist while generally being careful not to assume that rationality must be achievable when discussing his work. The notion of 'rationality' discussed below should therefore be read as broad enough to capture what Howson means by 'fairness'.

4.1.3   Subjective Bayesian inductive logic

Subjective Bayesian inductive logic is a research programme that aims to produce a framework for formalising the inductive assumptions that might be made by agents who are rational according to subjective Bayesian epistemology. It is distinct from subjective Bayesian epistemology in that it aims to produce formal models of inductive reasoning, rather than claims about the nature of rational partial belief. However, since it aims only to formalise the inductive reasoning of agents that subjective Bayesian epistemology would deem rational, it inherits the commitments of its epistemological counterpart.

Since subjective Bayesian inductive logic is committed to probabilistic necessity and sufficiency, it considers only probabilistic systems of inductive logic to be suitable as formal models of inductive reasoning. In addition, subjective Bayesian inductive logic follows subjective Bayesian epistemology's commitment to conditionalisation by stipulating that evidence should be represented using the 'conditional probability method'. According to this method, any possible evidence must be represented by a member of the domain of a probability function, and the effect of taking into account evidence represented by the domain-member E must be represented by differences between a pre-evidence probability function $pr_{initial}(\bullet)$ and a post-evidence function $pr_{new}(\bullet) = pr_{initial}(\bullet \mid E)$.

Subjective Bayesian inductive logic has been defended by Colin Howson in a series of published works, including Howson (1997b), Howson (1997a), Howson (2000), Howson (2001), Howson (2003), Howson (2009) and Howson (2012). Howson argues that subjective Bayesian inductive logic began with the work of Ramsey, citing passages such as the following:

> 'We find, therefore, that a precise account of the nature of partial belief reveals that the laws of probability are laws of consistency, an

> extension to partial beliefs of formal logic, the logic of consistency.'
> (Ramsey, 1931, p.182)

Romeijn expresses a similar point of view to Howson's in Romeijn (2005) and Romeijn (2011).

Subjective Bayesian inductive logic is similar to Carnapian inductive logic in that it represents inductive assumptions using axiomatically defined sets of measure functions with domains that are interpreted as representing possible objects of belief; that is, using systems of inductive logic. However, subjective Bayesian inductive logic is more restrictive than Carnapian inductive logic in several important ways.

1. **Focus on rational assumptions** First, it restricts itself to the task of formalising the inductive assumptions that might be made by rational agents, rather than addressing the problem of formalising inductive assumptions in general like Carnapian inductive logic.

2. **Special status of probability axioms** Second, whereas, according to Carnapian inductive logic, there is no distinction in principle between the probability axioms and any other formal axioms for measure functions, subjective Bayesian inductive logic endows the probability axioms with a special epistemological status.

3. **The conditional probability method** Finally, subjective Bayesian inductive logic is committed to representing evidence using the conditional probability method, rather than any other method. Carnapian inductive logic makes no such commitment, as can be seen, for example, from Carnap's remarks on imposing axioms based on informal evidential assessments at (Carnap, 1971a, p.120), and from his consideration of an alternative to the conditional probability method in Carnap (1967b).

In what follows I focus on the second of these differences between subjective Bayesian and Carnapian inductive logic, arguing that probabilistic necessity and sufficiency are epistemologically dubious. The main message of this chapter is therefore that subjective Bayesian inductive logic is epistemologically dubious. I believe that this issue shows decisively that Carnapian inductive logic is preferable. However, the other differences also weigh in favour of Carnapian inductive logic. In the case of the first difference, it seems arbitrary to limit inductive logic to representing assumptions that rational agents might entertain: a fully comprehensive inductive logic should surely go further. As for subjective Bayesian inductive logic's commitment to the conditional probability method of representing evidence, I argue below that it is methodologically dubious, though not at the same length as in the case of probabilistic necessity and sufficiency. The reader should therefore bear in mind that the epistemological problems highlighted below are not the only disadvantages of subjective Bayesian inductive logic.

Some terminology

In what follows I use the word 'probabilistic' as part of several artificial terms.

Given an inductive logical domain D, the symbol $\mathbb{P}_D$ stands for the set of probability functions on D. Where the domain is clear from the context I shall sometimes simply use the symbol $\mathbb{P}$. The axioms determining this set shall sometimes be called "probabilism".

I label as 'super-probabilistic' any collection of axioms that is strictly more restrictive than probabilism. In other words, given a domain D, super-probabilistic axioms determine sets of measure functions that are proper subsets of $\mathbb{P}_D$.

Conversely, 'sub-probabilistic' sets of axioms are strictly less restrictive than probabilism; they determine sets of measure functions of which $\mathbb{P}_D$ is a proper subset.

I call a system of inductive logic $(D, \mathcal{M})$ 'probabilistic' if $\mathcal{M} = \mathbb{P}_D$, 'superprobabilistic' if $\mathcal{M} \subset \mathbb{P}_D$ and 'sub-probabilistic' if $\mathcal{M} \supset \mathbb{P}_D$.

A state of belief is 'probabilistic' if it can be represented by a measure function of a probabilistic system of inductive logic.

4.2   Probabilistic necessity

This section evaluates subjective Bayesian inductive logic's commitment to probabilistic necessity. I argue that it has not yet been established that every possible rational state of partial belief can be represented by a probability function: probabilistic necessity is therefore a dubious commitment.

I consider three kinds of argument for probabilistic necessity. Betting arguments seek to show that the owners of non-probabilistic states of belief behave unwisely in a certain scenario. Axiomatic arguments seek to show that, if a state of belief cannot be represented by a probability function, then it necessarily violates a formal axiom that all rational states of partial belief ought to satisfy. Accuracy arguments claim that all non-probabilistic states of belief are unnecessarily inaccurate.

In their currently available forms, I argue, none of these kinds of argument establishes probabilistic necessity. Each kind of argument has unique problems. Betting arguments suffer from scenario-relativity and the packaging objection; currently available axiomatic arguments impose implausibly strict requirements on rational states of belief; accuracy arguments make unwarranted claims about the nature of inaccuracy. In addition, the arguments share several questionable assumptions and are similar in various respects: as a result they do not combine into a collectively compelling argument for probabilistic necessity.

4.2.1 Betting arguments

Betting arguments for probabilistic necessity seek to establish the irrationality of all non-probabilistic states of belief by exploiting a postulated link between states of belief and betting behaviour.

Every betting argument appeals to a hypothetical betting scenario which is supposed to makes this link apparent. In the scenario, it is claimed, agents with non-probabilistic states of belief have dispositions to make unwise bets. Such dispositions are taken to indicate irrationality.

The classic Dutch book argument proceeds as follows. First it is supposed that an agent's beliefs concern the set of sentences SL_{prop} of a propositional language L_{prop}, and that the true state of the world is described by a valuation function $v : SL_{prop} \rightarrow \{0, 1\}$ that assigns the value 1 to true sentences and 0 to false ones. Valuation functions must agree with the rules of classical propositional logic, so that, for example, if $v(\theta) = 1$ and $v(\phi) = 1$, then $v(\theta \wedge \phi) = 1$.[2]

Second, a scenario—the 'classic Dutch book scenario'—is envisaged where, given a set of sentences $B = \{\theta_1, \ldots \theta_n\} \subseteq SL_{prop}$, the agent is forced to select a 'betting' function $bet : B \rightarrow \mathbb{R}$, upon which choice an adversarial being will choose a 'stake' function $s : B \rightarrow \mathbb{R}$ and present the agent with an outcome, for every sentence $\theta \in B$, with value to the agent of $s(\theta)(v(\theta) - bet(\theta))$.

The classic Dutch book scenario is motivated by the idea that, for each sentence θ, the number $bet(\theta)$ represents odds at which the agent is prepared to accept either side of a bet on whether or not θ is true, with an arbitrary positive stake. The size of the stake is determined by $s(\theta)$'s absolute value, while this quantity's sign indicates which side of the bet the agent takes. Thus the agent's behaviour in the classic scenario arguably reflects how they are generally disposed to bet on the relevant sentences.

In order to connect behaviour in the classic Dutch book scenario with irrationality, a technical term—'bookability'—is defined. A betting function is 'bookable' if and only if there is a set of sentences $B^\star = \{\theta_1, \ldots \theta_n\} \subseteq B$ and a stake function $s^\star$ such that $\sum_{i=1}^{n} s^\star(\theta_i)(v(\theta_i) - bet(\theta_i))$ is negative for all possible valuation functions.

It is claimed that, if an agent chooses a bookable betting function in the classic Dutch book scenario, then they must have a pathological betting disposition. Such betting functions, it is argued, represent a willingness to accept combinations of bets that lose no matter what the true state of the world is.

Furthermore, it is supposed that such a pathological betting disposition indicates an irrational state of belief.

---

[2]Equivalently, the targets of belief could be supposed to be the members, or 'propositions', of a boolean algebra generated by an underlying state space, with the true state of the world represented by one member of the state space. I prefer to use logical terminology in order to emphasise common ground with the rest of my discussion.

Finally, betting functions are shown to be non-bookable in the classic Dutch book scenario if and only if they are probability functions. See Kemeny (1955) for a proof of the full biconditional. This is taken to show that agents whose states of belief cannot be represented by probability functions are irrational.

Objections to the classic Dutch book argument

This summary should have made clear that the classic Dutch book argument involves many potentially questionable assumptions. Every time one thing is taken to indicate another there is room to doubt the inevitability of the postulated link. Thus it is possible to doubt the links that the classic Dutch book argument postulates between having a non-probabilistic state of belief and choosing a non-probabilistic betting function, the link between choosing a bookable betting function and having a pathological betting disposition and the link between having a pathological betting disposition and having an irrational state of belief.

Below I focus on two objections to the link between choosing a bookable betting function in the classic Dutch book scenario and having a pathological betting disposition. According to these objections, it is sometimes possible for someone with a non-pathological betting disposition to choose a bookable betting function in the classic Dutch book scenario.

This kind of objection to the classic Dutch book argument is attractive, compared to other potential objections, because it involves fairly weak argumentative commitments. It is compatible with the view that bookability in the classic Dutch book argument is a useful proxy for irrationality in the vast majority of cases, and that irrationality always manifests itself as pathological betting behaviour of some kind. All that is required for the objection to succeed is that, in some cases, bookability in the classic Dutch book scenario can occur at the same time as a sensible, non-pathological, betting disposition.

The first objection compares the classic Dutch book scenario with another—the 'intuitionistic Dutch book scenario'—which is plausibly a more reliable guide to betting dispositions, yet does not yield probabilism as a rationality constraint. The second asserts that the link between bookability and pathologicality fails in cases where agents aggregate the values of combinations of bets by other means than addition.

I argue that both objections are successful and that, as a result, the classic Dutch book argument does not establish probabilistic necessity.

The intuitionistic Dutch book argument

There are other means that could be used to evaluate an agent's betting disposition, which yield results that conflict with the classic Dutch book argument.

For example, the Czech book argument, which features in Hájek (2008), seemingly shows that agents in a certain alternative scenario should choose a non-probabilistic betting function for exactly the same reason that an agent in the classic Dutch book scenario should choose a probabilistic one.

The question therefore arises as to why choice of a non-probabilistic betting function indicates a pathological betting disposition in the classic Dutch book scenario but not in alternative scenarios like the Czech Book scenario.

The natural response to this objection asserts that the classic Dutch book scenario is special because it corresponds to the kind of situations that occur in everyday life when people find themselves obliged to make decisions in the face of uncertainty. In the same way as situations that occur in sporting events arguably elicit athletes' true characters, despite being highly artificial, this response claims, the classic Dutch book scenario reveals agents' true betting dispositions. If it is the best scenario for this kind of elicitation, then this explains why the classic Dutch book scenario is special.

However, as I shall now argue, there is another scenario—the intuitionistic Dutch book scenario—which is even better than the classic Dutch book scenario at eliciting agents' betting dispositions, but conflicts with its conclusion.

The intuitionistic Dutch book scenario

The intuitionistic Dutch book scenario, set out at (Weatherson, 2003, §4), is very similar to the classic Dutch book scenario, but does not render all agents with non-probabilistic betting functions pathological. In this section I explain the intuitionistic Dutch book scenario, and argue that, in cases where the classic and intuitionistic Dutch book scenarios disagree, the latter elicits betting dispositions more reliably.

Like the classic Dutch book scenario, the intuitionistic Dutch book scenario envisages an agent whose beliefs concern the sentences of a propositional language. However, the intuitionistic scenario differs from the classic Dutch book scenario in the nature of its valuation functions, and in the way in which bets are settled.

Intuitionistic valuation functions, which I shall refer to with the symbol v_{int}, are identical to the valuation functions that appear in the classic Dutch book scenario apart from the following characteristics. Instead of being two-valued functions with the range $\{0, 1\}$, as in the classic Dutch book scenario, valuation functions in the intuitionistic Dutch book scenario have the range $\{0, \frac{1}{2}, 1\}$. Intuitionistic valuation functions assign the number 1 to sentences whose truth is revealed after a certain amount of time, 0 to sentences whose falsity is revealed and $\frac{1}{2}$ to all other sentences.[3] Instead of agreeing with the rules of classical logic,

---

[3]This account simplifies the standard account of intuitionistic semantics. In general, the range of an intuitionistic valuation function can be any Heyting algebra, allowing finer-grained distinctions between different un-decidable sentences: see (Palmgren, 2009, §2) for details. I restrict my attention to valuation functions with the range $\{0, \frac{1}{2}, 1\}$, as this is the simplest Heyting algebra that is not a Boolean algebra. The relevant points about the intuitionistic

intuitionistic valuation functions must agree with the rules of intuitionistic logic in order to be admissible.

Bets in the intuitionistic Dutch book scenario are settled as follows. Just as in the classic Dutch book scenario, a set $B \subseteq SL_{prop}$ is unveiled, the agent chooses a betting function $bet : B \to \mathbb{R}$ and a malevolent being chooses a stake function $s : B \to \mathbb{R}$. For any sentence θ, betting function bet and stake function s, if $v_{int}(\theta) \in \{0, 1\}$, then the value $val(bet, s, v_{int}, \theta)$ of the agent's outcome is $s(\theta)(v_{int}(\theta) - bet(\theta))$, just as in the classic case. If $v_{int}(\theta) = \frac{1}{2}$, then the value of the agent's outcome is $s(\theta)(0 - bet(\theta))$. In other words, if a sentence's truth is never revealed then the agent receives the same outcome as they would have received had it been revealed to be false.

Bookability in the intuitionistic Dutch book scenario is determined in essentially the same way as in the classic scenario. An agent with betting function bet is bookable in the intuitionistic scenario if and only if there is a set of sentences $B^\star \subseteq B$ and a stake function $s^\star$ such that $\sum_{\theta \in B^\star} val(bet, s^\star, v_{int}, theta)$ is less than zero for all possible intuitionistic valuation functions.

Weatherson proves that betting functions are un-bookable in the intuitionistic Dutch book scenario if and only if they can be extended into what he calls 'intuitionistic' probability functions, that is, measure functions satisfying the following collection of axioms:

> Intuitionistic probability axioms (propositional domains): $\mathbf{IP_{prop}}$
> Let $\vdash_{Int}$ represent the intuitionistic consequence relation. For any sentences θ and $\phi \in SL_{prop}$, $m : SL_{prop} \to \mathbb{R}$ must be such that:
>
> IP0 If $\theta \vdash_{Int} \phi$ for all $\phi \in SL_{prop}$, then $m(\theta) = 0$
>
> IP1 If $\top \vdash_{Int} \theta$, then $m(\theta) = 1$
>
> IP2 If $\theta \vdash_{Int} \phi$, then $m(\theta) \leq m(\phi)$
>
> IP3 $m(\theta) + m(\phi) = m(\theta \vee \phi) + m(\theta \wedge \phi)$

Relevance to the classic Dutch book argument's plausibility

Two key features of the intuitionistic Dutch book argument make it relevant to the classic Dutch book argument's plausibility. First, not all intuitionistic probability functions are classical probability functions - as a result, a betting function can be bookable in the classic Dutch book scenario but un-bookable in the intuitionistic scenario. Second, the intuitionistic scenario is a more realistic tool for eliciting betting dispositions than the classic Dutch book scenario.

The first feature can be demonstrated as follows. First note that substituting the classical consequence relation $\vdash_{Cl}$ for the intuitionistic one $\vdash_{Int}$ in $\mathbf{IP_{prop}}$ makes these axioms equivalent to probabilism. This is shown at (Weatherson, 2003, p.2). Next, consider that, whenever a sentence intuitionistically entails

---

scenario are already apparent in this simple case and extend to more complicated ones.

another, it also classically entails it, but that the converse relationship does not hold. For example, $\top$ entails $(\theta \vee \neg\theta)$ classically but not intuitionistically.

As a result, for each of the conditions in $IP0 - IP2$ there are cases where the antecedent is satisfied on the classical interpretation of the consequence relation but not on the intuitionistic interpretation. In these cases, classical probability functions have to conform to restrictions that do not apply to intuitionistic probability functions. The reverse is never the case, and $IP3$ applies to both interpretations equally. The requirement to be a classical probability function is therefore more restrictive than the requirement to be an intuitionistic probability function.

As a result of this fact, it is possible that a betting function *bet* is an intuitionistic probability function but not a classical probability function. In this case, *bet* would be bookable in the classic dutch book scenario but not in the intuitionistic Dutch book scenario.

The intuitionistic Dutch book argument's superiority as a means for eliciting betting dispositions is shown by an argument in Harman (1983). When people make bets in real life, they often face the possibility that the truth or falsity of the propositions they bet on will not be revealed in time for the bet to be settled. For example, when betting on whether or not a certain government minister called a policeman a 'pleb' during an altercation, I must take into account the fact that, whatever the truth or falsity or this proposition, it may never be possible to settle a bet on the matter. The key question is not the proposition's truth, but whether or not it has been revealed to be true. Harman argues that this kind of situation is typical of real-life betting.

The intuitionistic Dutch book scenario accommodates this possibility without difficulty because, as we saw above, intuitionistic valuation functions v_{int} : $B \rightarrow \{0, \frac{1}{2}, 1\}$ can naturally be interpreted as representing whether a proposition's truth has been revealed. On the other hand, classical valuation functions $v : SL_{prop} \rightarrow \{0,1\}$ cannot easily be interpreted in this way: consequently the classic Dutch book scenario must stipulate that the truth or falsity of all propositions is revealed before prizes are allocated. This unrealistic stipulation makes the classic Dutch book scenario a less reliable guide to betting dispositions than the intuitionistic Dutch book scenario.

Since the intuitionistic Dutch book scenario is a more realistic guide to betting dispositions than the classic Dutch book scenario, its bookability constraints must be better indicators of irrationality than those of the classic scenario. However, the intuitionistic Dutch book argument is not restrictive enough to ensure that all non-classically-probabilistic betting functions are bookable. Making the classic Dutch book argument more realistic therefore seems to turn it from an argument for probabilistic necessity to an argument against it. This fact undermines betting arguments for probabilistic necessity.

The packaging objection

A second objection to betting arguments for probabilistic necessity—the 'packaging objection'—asserts that the classic Dutch book argument's definition of bookability does not reflect betting in real life.

According to the classic Dutch book argument's definition, a betting function is bookable if there is a set of sentences and a stake function such that the sum of the values of the sentences' associated outcomes is negative, whatever the true valuation function. The choice of a bookable betting function is taken to indicate a pathological betting disposition, because it seemingly represents willingness to accept a combination of bets that has a negative aggregate value in all circumstances.

However, this line of reasoning tacitly assumes that, at least in cases relevant to the classic Dutch book argument, the aggregate value of a combinations of bet-outcomes is found by adding together the values of the individual outcomes. The assumption that the aggregate value of a combination or 'package' of bet-outcomes can always be found in this way, known as the 'package principle', is what the packaging objection contests.

Schick argues at (Schick, 1986, p.114) that the package principle amounts to an assumption of 'value-wise independence', according to which the impact of adding a new outcome to a combination is the same, as far as the value of the new combination is concerned, whatever the nature of the initial combination. Schick claims that value-wise independence typically does not obtain in real-life betting situations: adding a new outcome often has a different effect depeding on the nature of the initial package.

The following example illustrates Schick's point: a mildly bad bet outcome—say, losing one pound out of a budget of fifty—might seem relatively innocuous when considered individually. However, if it were part of a potential combination of bet-outcomes that would leave one's overall wealth below a certain threshold, such as the cost of a bus ride home from the bookmaker's shop, the value of the outcome might seem different.

Similarly, if one needs to gain a certain amount of money, for example to pay a ransom, one might take a bet in combination with other bets that would seem imprudent in isolation. If a bet's outcomes are not monetary, then value-wise independence becomes even harder to achieve: for example, the prospect of winning a sausage will surely add more value to a package of bets that may result in the acquisition of potatoes than it will to a package involving lemons.

Schick's objection demonstrates that the package principle does not always obtain in real-life betting situations. As a result, it is not always safe to assume that an agent's choice of a bookable betting function in the classic Dutch book scenario shows that, in the real world, they are willing to accept a combination of bets that has negative aggregate value, whatever the state of the world. It could

be that some agents who choose bookable betting functions do not have pathological betting dispositions, but merely aggregate the values of bet-outcomes using means other than summation.

The classic Dutch book argument could perhaps be saved from the packaging objection by imposing value-wise independence on the Dutch book scenario by stipulation. This response requires a corresponding restriction of the Dutch book argument's scope: instead of throwing light on agents' betting dispositions in general, the classic Dutch book scenario must be taken as showing nothing beyond their dispositions to bet in situations where value-wise independence obtains.

However, there are very few real life betting situations where value-wise independence obtains. Real bettors typically have limited budgets, and are therefore potentially susceptible to avoiding bets when they are part of potentially-bankrupting packages. On the other hand, in situations like the ransom case above, it can also be sensible to avoid a bet on the grounds that the package it appears in is too small. Such situations are plausibly fairly common: for example, one might not be inclined to make any bets at all unless a package of bets is available with the potential to produce enough money to outweigh the cost in time and effort of going to a bookmaker's shop.

Since value-dependence is so ubiquitous in real-life betting situations, imposing value-independence on the Dutch book scenario seems to make it unrealistic and is therefore not a viable response.

Conclusion

The intutionistic Dutch book objection and the packaging objection show that betting arguments cannot currently establish probabilistic necessity. Both objections show that agents who choose bookable betting functions in the classic Dutch book scenario do not necessarily have pathological betting dispositions. Consequently, betting arguments cannot show that all non-probabilistic states of belief are irrational.

4.2.2 Cox's axiomatic argument

Axiomatic arguments aim to identify collections of requirements that, it is claimed, measure functions must satisfy in order to represent rational states of partial belief. They aim to demonstrate that all such measure functions are probability functions.

Together with preliminary assumptions guaranteeing that all rational states of partial belief can be represented by measure functions, axiomatic arguments can be used to argue for probabilistic necessity.

The most famous axiomatic argument is Cox's, found in Cox (1947), which I focus on below. I reproduce Cox's argument in a somewhat modified form which I hope will be more easily understood than Cox's in the context of the preceding discussion.

I argue that Cox's argument is unsuccessful because some of its requirements are not genuine rationality constraints.

Cox's argument

Preliminary assumptions
Like the classic Dutch book argument, Cox's argument imposes a preliminary 'classicality' requirement on rational partial beliefs, though it is slightly different in this case. In order to satisfy Cox's preliminary requirement, a state of partial belief must concern the sentences SL_{prop}^{∞} of a propositional language L_{prop}^{∞} with infinitely many propositional variables. In addition, a classically admissible valuation function $v : SL_{prop}^{\infty} \to \{0, 1\}$ must identify which of the sentences in SL_{prop}^{∞} are true and false.

A second preliminary assumption of Cox's argument is that every state of belief is representable by a 'belief function' with the form $b : SL_{prop}^{\infty} \to [0, 1]$. 'Conditional' belief functions are defined in the usual way, so that $b_{cond}(\theta \mid \psi) = \frac{b(\theta \wedge \psi)}{b(\psi)}$ if $b(\psi) \neq 0$ and are undefined otherwise. This assumption implies that rational states of partial belief must be 'sharp', that is, representable by single real numbers. This rules out alternative approaches which might use other mathematical objects to represent rational partial belief, such as pairs of real numbers or intervals.

Axioms
Having made these two preliminary assumptions, which ensure that rational states of belief can be represented by belief functions, Cox's argument imposes additional rationality requirements in the form of axioms for belief functions.

It is not clear that the axioms in Cox's original presentation are sufficient to prove the mathematical result that underpins his argument: see Halpern (1999a) for details on this point. I therefore present below the less controversial axioms from (Paris, 1994, p.24), which do not suffer from this problem.

> **Cox's axioms**
> For all sentences θ, ψ, $\phi \in SL_{prop}^{\infty}$, if $b : SL_{prop}^{\infty} \to [0, 1]$ represents a rational state of belief, then:
>
> C1 If $\vDash \theta \iff \theta'$ and $\vDash \psi \iff \psi'$, then $b_{cond}(\theta \mid \psi) = b_{cond}(\theta' \mid \psi')$.
>
> C2 If $\vDash \psi \implies \theta$, then $b_{cond}(\theta \mid \psi) = 1$ and $b_{cond}(\neg\theta \mid \psi) = 0$.

C3 There exists a continuous function $F : [0,1] \times [0,1] \to [0,1]$ which is strictly increasing on $(0,1] \times (0,1]$ and such that $b_{cond}(\theta \wedge \phi \mid \psi) = F(b_{cond}(\theta \mid \phi \wedge \psi), b_{cond}(\phi \mid \psi))$.

C4 There exists a decreasing function $S : [0,1] \to [0,1]$ such that $b_{cond}(\neg\theta \mid \psi) = S\left(b_{cond}\left(\theta \mid \psi\right)\right)$.

C5 For any real numbers $\epsilon > 0$ and $0 \leq \alpha, \beta, \gamma \leq 1$ there are sentences $\theta_1, \theta_2, \theta_3$ and θ_4 such that:

- It is not the case that $\vDash \neg(\theta_1 \wedge \theta_2 \wedge \theta_3)$

- $\mid b_{cond}(\theta_4 \mid \theta_1 \wedge \theta_2 \wedge \theta_3) - \alpha \mid < \epsilon$.

- $\mid b_{cond}(\theta_3 \mid \theta_1 \wedge \theta_2) - \beta \mid < \epsilon$.

- $\mid b_{cond}(\theta_2 \mid \theta_1) - \gamma \mid < \epsilon$.

C1 and C2 impose restrictions on conditional belief functions based on deductive logic, forcing them to behave as expected in cases of logical equivalence and implication.

C3 and C4 impose functional relationships on conditional belief functions based on intuitions about rational conditional belief in conjunctions and negations respectively.

C3 requires that belief in any conjunction $\theta \wedge \phi$, conditional on any other sentence ψ, is determined by the pair $(b(\theta \mid \phi \wedge \psi), b(\phi \mid \psi))$ consisting of the conditional belief in θ given $\phi \wedge \psi$ and the conditional belief in ϕ given ψ. In addition, the function F expressing this dependence must be continuous and, provided that both of these quantities are positive, strictly increasing.

Similarly, according to C4, belief in any negation $\neg\theta$, conditional on any sentence ψ, is determined by the value of $b(\theta \mid \psi)$, and the function S expressing this dependence must be decreasing, though it need not be continuous. This requirement ensures that, the more ψ increases θ's credibility, the more it decreases the credibility of $\neg\theta$, and is motivated by the intuition that rational beliefs in negations behave in this way.

C5 requires that, given any three 'target' numbers α, β and γ between zero and one and any rational belief function b_{cond}, it is possible to choose sentences θ_1, θ_2, θ_3 and θ_4 in such a way that the specified conditional beliefs take values that are arbitrarily close to their targets.

Cox's theorem

The centrepiece of Cox's argument is Cox's theorem.

Cox's Theorem For any belief function $b : SL^{\infty}_{prop} \to [0,1]$ that satisfies C1-C5, there is a continuous, strictly increasing, surjective function $g : [0,1] \to [0,1]$ such that $g \circ b : SL^{\infty}_{prop} \to [0,1]$ is a probability function and $g \circ b(\theta \wedge \psi) = g \circ b(\theta \mid \psi) \times g \circ b(\psi)$.

This result is proved at (Paris, 1994, p.25-32).

Cox's theorem shows that any state of belief satisfying classicality, sharpness and C1-C5 can be represented by a belief function that is also a probability function. Consequently, if all of these requirements were genuine rationality constraints, then Cox's argument would underpin a sound epistemological argument for probabilistic necessity. However, as I shall now argue, this is not the case.

Criticism of Cox's requirements

Colyvan (2004) argues that the preliminary classicality requirement is unwarranted because it forces belief functions to conform to the law of excluded middle, which he argues does not always apply. In many situations, he claims, it is unsafe to assume that every proposition either is or is not true. Such danger may arise in connection with fictional discourse, vague properties and un-proved mathematical conjectures.

According to the sharpness requirement, rational states of partial belief should be represented by real-valued functions. Cox's argument therefore has to assume a difficult position regarding phenomena that are difficult for such functions to represent, such as higher-order uncertainty and beliefs with degrees that are impossible in principle to compare. It must either assert that these phenomena can safely be disregarded by formal rationality principles, or else that they can be represented in a sharp framework despite the apparent difficulty. See Dubois et al. (1996) for a discussion of why both of these positions are unappealing.

C1 and C2 are implausible in the absence of logical omniscience, as, in order to be sure that their belief function does not contravene these requirements, an agent would need to know all of the logical equivalences and implications of the language relevant to their beliefs.

Cox's argument shares these problematic assumptions with other arguments for probabilistic necessity: I discuss logical omniscience, sharpness and classicality below in section 4.2.4. Specific objections to Cox's argument arise in connection with axioms C3 and C5.

Against C3 it can be objected that the dependence this axiom requires is incorrect—that some values of $b_{cond}(\theta \wedge \phi \mid \psi)$ should not be determined by $b_{cond}(\theta \mid \phi \wedge \psi)$ and $b_{cond}(\phi \mid \psi)$—or that the function expressing the dependence should either be non-continuous or not strictly increasing. (Van Horn, 2003, §9) contains a discussion of the first point, arguing that independent influence by $b_{cond}(\theta \mid \psi)$ or $b_{cond}(\phi \mid \theta \wedge \psi)$, or both of these quantities, should not necessarily be ruled out on pain of irrationality.

The continuity of the function F is difficult to motivate on strictly epistemological grounds. It is not clear why it should be irrational for small changes

in $b_{cond}(\theta \mid \phi \wedge \psi)$ and $b_{cond}(\phi \mid \psi)$ to be associated with large changes in $b_{cond}(\theta \wedge \phi \mid \psi)$.

Finally, axiom C5 is also problematic. As observed at (Halpern, 1999a, p.69), C5 entails that all rational belief functions' ranges are dense subsets of $[0, 1]$. This is an undesirable property in itself, as it imposes irrationality on all believers whose degrees of belief are not boundlessly fine-grained. It is not clear that real states of partial belief are capable of such discernment. Whether or not it is plausible for real agents to be able to make arbitrarily fine-grained distinctions in their degrees of belief, this ability does not seem intuitively to be connected with rationality. Since any agent's fine-grainedness can only feasibly be measured so as to fix it within a given range, it will be impossible to determine whether any agent's beliefs are sufficiently fine-grained enough for them to satisfy this requirement. In addition, Halpern points out that $C5$ prevents Cox's theorem from applying to belief functions with finite domains. Again, it does not seem like whether or not an agent is rational should depends on how many propositions they are able to form beliefs about.

Conclusion

Cox's axioms C3 and C5 are not convincing epistemologically. For this reason, Cox's axiomatic argument does not provide convincing grounds on which to believe in probabilistic necessity. However, these axioms may be convincing methodologically, as approximations of rationality constraints that are 'good enough' for practical purposes. Van Horn (2003) offers this kind of methodological defences of C3 (§9) and C5 (§7), while (Halpern, 1999b, p.434) presents a version of Cox's argument where C5 is altered to allow belief functions to have finite domains. Halpern acknowledges that this modification involves sacrificing some epistemological plausibility, but defends his move on methodological grounds.

4.2.3 Accuracy arguments

Like axiomatic arguments, accuracy arguments consist of putative requirements of rational states of belief, namely that they be sharp, classical, respect logical equivalence and, most importantly, obey some decision-theoretic norm involving accuracy. However, accuracy arguments differ from axiomatic arguments because they also make descriptive claims about the nature of inaccuracy. Accuracy arguments are exemplified by Joyce (1998), Joyce (2009) and Leitgeb and Pettigrew (2010). Predd et al. (2009) contains a mathematical argument that can be used as the basis for an accuracy argument, though the authors do not explicitly endorse this approach.

Below I discuss accuracy arguments' preliminary claims, before focusing on some

specific claims about the nature of inaccuracy. I conclude that currently available accuracy arguments are open to objections at both levels.

Preliminary claims

Accuracy arguments make preliminary claims about states of belief, inaccuracy and rationality, which are meant to be relatively uncontroversial, setting the scene for more controversial subsequent claims about inaccuracy. I shall now discuss some of these.

States of belief

Like Cox's axiomatic argument, accuracy arguments begin with preliminary assumptions which ensure that states of belief that need to be evaluated for their accuracy can be represented by real-valued functions on an appropriate domain. These differ slightly from Cox's preliminary assumptions: whereas Cox's axioms jointly ensure that belief functions' domains are infinite, accuracy arguments typically assume that belief functions' domains are finite.

'Classicality' stipulates that states of belief concern the sentences SL_{prop} of a propositional language L_{prop} with finitely many propositional variables $A_1, \ldots, A_n$ and that the true state of the world is described by a classically admissible valuation function $v : SL_{prop} \to \{0, 1\}$. I label the set of all such valuation functions $\mathbb{V}$.

'Sharpness' forces states of belief to be representable by real-valued belief functions $b : SL_{prop} \to \mathbb{R}_{\geq 0}$.

'Respect for logical equivalence' forces belief functions to assign the same value to all logically equivalent sentences. For any sentences θ and ϕ, if $\theta \iff \phi$ is logically true, then it must be the case that $b(\theta) = b(\phi)$. I denote by $\{\theta\}^{\equiv}$ the equivalence class of sentences logically equivalent to θ and by $SL^{\equiv}_{prop}$ the set of all the equivalence classes of SL_{prop}. I typically use the symbol X to represent members of $SL^{\equiv}_{prop}$. $\mathbb{B}$ is the set of all belief functions satisfying respect for logical equivalence.

Equivalently, these assumptions can be formulated as the adoption of a framework in which beliefs concern the members or 'propositions' of the powerset of a state-space Ω with finitely many atomic states. In this case each atomic state $\omega \in \Omega$ represents a 'possible world' and propositions are said to be 'true at a particular world' if they contain that world's atomic state. The atomic states in this possible-world framework correspond exactly to the atomic sentences $\pm A_1 \wedge \ldots \wedge \pm A_n$ in SL_{prop}, and the propositions to the members of $SL^{\equiv}_{prop}$.

I discuss objections to these claims about states of partial belief below in section 4.2.4 as accuracy arguments share them with other arguments for probabilistic necessity.

Inaccuracy sharpness

Accuracy arguments make a preliminary sharpness assumption about inaccuracy, according to which the inaccuracies of states of belief can be measured by functions with the form $\mathbf{S} : \mathbb{B} \times \mathbb{V} \to \mathbb{R}_{\geq 0}$.

Each such 'inaccuracy measure' assigns a positive real number to every possible belief-function-valuation-function pair (b, v) according to how accurate the state of belief described by b is in the state of affairs designated by v. Following Leitgeb and Pettigrew (2010), I use the term 'inaccuracy measure' throughout, but the reader should be aware that terminology in the field of accuracy arguments has not yet settled, and 'epistemic scoring rule' and 'epistemic dis-utility function' are sometimes used elsewhere.

It is not uncontroversially true that every state of belief is inaccurate to a precise degree, as in some cases it seems that the inaccuracies of different states of beliefs are not directly comparable.

Consider, for example, two agents whose states of belief are reasonable and identical, except that agent A strongly believes that humans coexisted with dinosaurs, while agent B equally strongly believes that cheddar cheese tastes like cucumber. Due to the very different ways in which these beliefs are incorrect, it does not seem quite right to say that both agents' states of belief are equally inaccurate, nor that one is less accurate than the other. Instead, it is arguably more appropriate to describe the situation as one in which the two states of belief fail in different, incommensurable ways: agent A's state of belief is historically inaccurate, whereas agent B is guilty of gastronomic inaccuracy.

However, inaccuracy sharpness rules out this way of describing the situation, as it forces inaccuracy to be measurable along just one dimension. Given the varied ways in which states of belief can fail, insisting that overall inaccuracy can be measured by a single quantity seems unwarranted.

Accuracy-based epistemic norm

Having made these preliminary assumptions about states of belief and their inaccuracy, accuracy arguments can formulate their central rationality requirement as an epistemic norm identifying states of partial belief represented by certain belief functions as irrational based on their inaccuracy.

The epistemic norm that is most often invoked by accuracy arguments for probabilism is called 'dominance-avoidance'. A belief function b is said to 'accuracy-dominate' another belief function b' according to the inaccuracy measure $\mathbf{S}$ if and only if $\mathbf{S}(b, v) \leq \mathbf{S}(b', v)$ for all valuation functions $v \in \mathbb{V}$, with strict inequality holding for at least one valuation function.[4] According to the norm of dominance avoidance, accuracy-dominated belief functions represent irrational states of belief.

---

[4]It is assumed that all the functions mentioned here have a common domain.

To see why dominance-avoidance is seen as compelling, suppose that one belief function b_A representing agent A's state of partial belief dominates agent B's belief function b_B according to the inaccuracy measure S. Then, according to S, agent B is at least as inaccurate as agent A in all possible states of the world, and more inaccurate in at least one state. A's state of belief is more accurate than B's in this sense. Moreover, B could have achieved more accuracy without having to acquire any new evidence, simply by pondering the situation's formal representation. If B fails to do so, one might ask questions about their rationality.

I can think of two objections to dominance avoidance.

First, it is conceivable that situations might arise where the true state of the world is known to depend on an agent's state of belief. Such cases are analogous to act-dependence in decision theory, which is known to frustrate dominance reasoning. It is argued at Caie (2013) that the possibility of belief-dependence makes accuracy-dominance an unreliable norm. Greaves (2013) also points out cases where belief-dependence can complicate accuracy-based reasoning.

Secondly, a rational agent might choose an accuracy-dominated state of belief because the epistemic good of having an accurate state of belief is outweighed by another epistemic good, such as epistemic helpfulness.

For example, consider the case of Sherlock Holmes and Dr. Watson. In Arthur Conan Doyle's stories, these two agents form an epistemic community who collectively seek to solve mysteries. Watson has the knack of making observations and suggestions that, despite or even perhaps due to their foolishness, cause Holmes to make crucial deductions. In other words, Watson has a tendency to form inaccurate states of belief that are nonetheless epistemically helpful. If Watson were offered the opportunity to exchange his state of belief for a slightly more accurate but less helpful alternative—say by attending a criminology night-class—he might rationally decline. The increased accuracy is unlikely to make Watson's community able to solve any mysteries that it would not have solved anyway as Watson would merely learn things that Holmes knows already. However, the loss of an evening's worth of interactions might well cause some mysteries to go unsolved because of Holmes being unable to take advantage of Watson's capacity as an intellectual springboard.

Assessment of preliminary claims

Accuracy arguments aim to make comparatively uncontroversial preliminary claims about states of partial belief, the correct functional form of inaccuracy measures and the relationship between inaccuracy and irrationality.

The idea is to shift any controversy away from this kind of claim and towards claims about which inaccuracy measures are legitimate; that is, which measures genuinely represent inaccuracy. By proposing candidate legitimacy conditions, an argument for probabilism can be built up piece by piece. Eventually, it

is hoped, legitimacy conditions will be found that are both intuitively well-motivated and collectively powerful enough to ensure that all non-probabilistic belief functions are accuracy-dominated.

However, many preliminary assumptions made by accuracy arguments for probabilism are not as uncontroversial as intended. We have seen that it is not safe to assume that every state of belief is inaccurate to a real-valued degree, nor that all states of belief represented by accuracy-dominated belief functions are irrational. We shall see later that some preliminary assumptions about states of partial belief are also difficult to justify. There therefore seem to be problems with the overall strategy underlying accuracy arguments for probabilism.

Nonetheless, it will still be useful to consider the merits of some legitimacy conditions for inaccuracy measures that have been proposed recently. The next few subsections cover four of these: sum-decomposability, proposition-neutrality, strict propriety and continuity.

If a sufficiently compelling and strong collection of legitimacy conditions were found, then accuracy arguments for probabilism would apply in certain circumstances. Specifically, probabilism would be vindicated in circumstances where inaccuracy uncontroversially takes real-valued degrees and dominance-based reasoning is clearly appropriate.

On the other hand, if, as I shall argue, the sufficiently strong collections of conditions that have been proposed so far are dubiously compelling, then this would strengthen the case against accuracy arguments. As well as an external critique according to which accuracy arguments for probabilism have questionable presuppositions, there would also be an internal critique, according to which accuracy arguments are unsuccessful even when their presuppositions are granted.

Sum-decomposability

An important legitimacy condition for inaccuracy measures, made in most accuracy arguments that are currently thought to be viable, is sum-decomposability.

An inaccuracy measure $\mathbf{S} : \mathbb{B} \times \mathbb{V} \to \mathbb{R}$ satisfies sum-decomposability if it is related to a proposition-specific inaccuracy measure $\mathbf{L} : \Omega \times [0, 1] \times \{0, 1\} \to [0, 1]$ in such a way that the following equality holds in all cases:

$$\mathbf{S}(b, v) = \sum_{X \in \Omega} \mathbf{L}(X, b(X), v(X)). \tag{4.1}$$

Sum-decomposability ensures that the number representing the 'global' inaccuracy of a state of belief is the sum of the numbers representing the 'local' inaccuracies of its individual partial beliefs in particular propositions.

Leitgeb and Pettigrew argue at (Leitgeb and Pettigrew, 2010, §5.2.1) for sum-decomposability on the grounds that local, proposition-specific inaccuracy and

global, state-of-belief-oriented inaccuracy ought to be related in some systematic way. If there were no such systematic relationship, they claim, then the two conflicting forms of inaccuracy could force rational agents into dilemmas. The imperative to pick a state of belief with minimal inaccuracy would conflict with the imperative to choose individual beliefs each of which is minimally inaccurate. Legitimate inaccuracy measures, they argue, could never produce such situations. Accepting on these grounds that there must some systematic relationship between global and local inaccuracy Leitgeb and Pettigrew settle on summation as an easily available and well-understood candidate.

Against this line of argument it can be objected that, where the two conflict, global inaccuracy should always trump proposition-specific inaccuracy because of the need to evaluate features of states of belief that are irreducibly global, such as patterns between different beliefs.

Only global inaccuracy can evaluate states of belief in a holistic way, taking into account all degrees of beliefs in all propositions at the same time. Proposition-specific inaccuracy cannot do this and should therefore, I claim, be over-ruled in cases of disagreement.

To illustrate this point, suppose that agent A is particularly inaccurate when it comes to propositions about their favourite football team. Whereas the inaccuracies of their beliefs about other topics are unremarkable, they inaccurately believe that the red team will win all of its matches, that it plays an intrinsically more beautiful game, that its players are excellent people, and so on. Intuitively, this pattern in agent A's beliefs should be relevant to the inaccuracy of their overall state of belief. Arguably the agent's overall state of belief is less inaccurate than that of agent B whose equal number of equally inaccurate beliefs are distributed randomly among many different topics. Alternatively, agent A's clustered inaccurate beliefs might be thought to make them *more* globally inaccurate than agent B's evenly spread ones. The author's intuitions on this point are not entirely conclusive: however it seems clear is that it is not safe to assume that A and B have the same global inaccuracy.

Whether agent A's clustering brings about an extra penalty or a partial reprieve, the two agents would have identical local inaccuracies but different global inaccuracies. Faced with this conflict, it seems clear that we should evaluate the two agents according to their global inaccuracy rather than their local inaccuracy. To do the reverse would neglect the importance of the pattern in agent A's state of belief.

However, since sum-decomposability imposes a systematic relationship between local and global inaccuracy measures, it prevents global inaccuracy from trumping local inaccuracy in this way. It therefore ignores an important aspect of inaccuracy.

Another problem with sum-decomposability concerns the nature of the relationship between local and global inaccuracy, supposing that one exists. As Landes (2015) argues and Leitgeb and Pettigrew admit, there is no particu-

lar reason why this relationship should be described by summation over all propositions rather than by any other operation. The conflict between global and proposition-specific imperatives could have been avoided by an alternative form of decomposability, such as one which required global inaccuracy to be the product of proposition-specific inaccuracies rather than their sum. Leitgeb and Pettigrew's argument does not, therefore, go far enough to fully justify sum-decomposability, but only a weaker restriction asserting that legitimate inaccuracy measures should be decomposable in some way or another. It is not yet clear whether a successful accuracy argument can proceed on the basis of such a weaker form of decomposability.

To sum up, it is not clear that there is a systematic relationship between global and local inaccuracy. Even if there is such a relationship, a compelling case that it is described by summation over all propositions has not yet been made.

Proposition Neutrality

Proposition neutrality is a condition on local, proposition-specific, inaccuracy measures. It stipulates that, for any legitimate local inaccuracy measure $\mathbf{L}$, if $b(X) = b(Y)$ and $v(X) = v(Y)$, then $\mathbf{L}(X, b(X), v(X)) = \mathbf{L}(Y, b(Y), v(Y))$. Proposition neutrality ensures that, if two propositions have the same truth values and are believed to the same degree, then they each contribute the same amount of local inaccuracy. If a local inaccuracy measure is proposition-neutral, then it is effectively a function with only two arguments—a degree of belief and a truth value—any influence by the first argument is ruled out.

Many prominent accuracy arguments for probabilistic necessity, including Pettigrew (ming) and Leitgeb and Pettigrew (2010) suppose that legitimate local inaccuracy measures are proposition-neutral.

However, proposition-neutrality is an unattractive legitimacy condition because it prevents local inaccuracy from taking into account evidence. Whether or not a proposition is supported by one's evidence, believing it strongly will contribute the same amount to one's inaccuracy, according to a proposition-neutral local inaccuracy measure.

The following example illustrates this point. Suppose that agents A and B are each presented with evidence supporting the proposition that a certain patient has cancer: the result of a very reliable test was positive. Having taken this evidence into account, agent A believes this proposition strongly, whereas agent B ignores the evidence, taking the more ambivalent position of believing the proposition to the degree $1/2$. In addition, both agents have attitudes towards the proposition that the number of fish on Earth is odd, about which they have minimal evidence. Agent A believes this second proposition to degree $1/2$, whereas agent B believes it very strongly: in fact exactly as strongly as agent A believes that the patient has cancer. Suppose that both propositions happen to have the same truth value: either the patient has cancer and there are an odd

number of fish on Earth, or the patient doesn't have cancer and the number of fish is even.

Agent A's beliefs seem to be appropriate in view of the available evidence, whereas agent B's beliefs are very inappropriate. In this sense agent A seems to be more accurate than agent B. However, this difference cannot easily be captured by a proposition-neutral inaccuracy measure.

To see why, suppose that the two propositions are represented by members Can and Odd of an appropriate set Ω, that the belief function b_A represents agent A's beliefs and that b_B represents those of agent B. Suppose that the local inaccuracies of b_A and b_b are measured by a proposition-neutral local inaccuracy **L**

Since **L** is proposition neutral, and given the stipulated facts about the propositions' truth values and the agents' degrees of belief, the following equalities will hold:

- $\mathbf{L}(Can, b_A(Can), v(Can)) = \mathbf{L}(Odd, b_B(Odd), v(Odd))$

- $\mathbf{L}(Odd, b_A(Odd), v(Odd)) = \mathbf{L}(Can, b_B(Can), v(Can))$

Despite the evidential mismatch, the local inaccuracy of agent A's belief in Can will be represented as being the same as that of agent B's equally strong belief in Odd, and the local inaccuracy of agent A's belief in Odd will be the same as that of agent B's equally strong belief in Can. As a result, contrary to the intuition that agent A's beliefs are more accurate than agent B's, **L** will give b_A and b_B the same total scores with respect to Can and Odd.

Two possible responses to this objection proceed as follows.

First, a proponent of proposition-neutrality might reject the intuition according to which agent A's beliefs are more accurate than agent B's, arguing that the two agents are actually equally inaccurate. This would seem to imply that one can sometimes avoid inaccuracy while responding inappropriately to one's evidence. This would be puzzling as taking evidence into account appropriately seems to be a key component of rationality.

This response therefore weakens the key claim of accuracy arguments that all accuracy-dominated belief functions represent irrational states of partial beliefs. If accuracy and responding appropriately to evidence can come apart in this way, then some accuracy-dominated belief functions may represent rational states of partial belief that take evidence into account sensibly. Similarly, some non-dominated belief functions might represent irrational states of belief that fail to do so.

Secondly, the proponent of proposition-neutrality might argue that a good global inaccuracy measure would give b_A a lower overall aggregate score than b_B in the case described above, despite the fact that both belief functions would have the same local scores with respect to Can and Odd. There would have to be a difference in b_A and b_B with respect to other members of Ω. In order for this

response to be convincing, however, it would need to be shown with respect to which members of Ω it would be safe to assume that b_B is more inaccurate than b_A.

Neither of these responses seems promising. As a result the status of proposition-neutrality as a genuine legitimate condition for local inaccuracy measures is dubious.

Strict propriety

Strict propriety is another condition on local inaccuracy measures. According to (Landes, 2015, p.9)[5], a local inaccuracy measure $\mathbf{L}$ is strictly proper if and only if it satisfies the following two properties:

- The expression $\mathbf{L}(\Omega, b(\Omega), 1) + \mathbf{L}(\emptyset, b(\emptyset), 0)$ is uniquely minimised when $b(\Omega) = 1$ and $b(\emptyset) = 0$.

- If X is not Ω or $\emptyset$, then for all $p \in [0, 1]$, the expression $p\mathbf{L}(X, b(X), 1) + (1 - p)\mathbf{L}(X, b(X), 0)$ is uniquely minimised when $b(X) = p$.

Strict propriety has received a lot of attention in the literature on accuracy arguments - see, for example, (Joyce, 2009, §8), Pettigrew (2013) and (Pettigrew, ming, Ch. 4). Leitgeb and Pettigrew (2010) restrict their attention to quadratic inaccuracy measures that are strictly proper, but do so not because quadratic measures have this property, but rather because they do not place agents in epistemic dilemmas. However, the dilemmas that Leitgeb and Pettigrew envisage involve imperatives to minimise different forms of *expected* inacccuracy, and the meat of my objection to strict propriety is that non-probabilistic agents need not obey such imperatives. Consequently, my objection may also pose problems for Leitgeb and Pettigrew's accuracy argument.

Pettigrew hints at a possible motivation for the conclusion that legitimate inaccuracy measures are strictly proper, according to his definition of this notion[6]:

> [Provided that] a scoring rule is proper, then, a probabilistic agent with credence p in a proposition expects that credence and only that credence to be epistemically best.
> (Pettigrew, 2013, p.28)

Our definition of strict propriety validates this observation. Suppose that $\mathbf{L}$ is a strictly proper local inaccuracy measure, and an agent has the probabilistic credence p in a proposition X. The expected proposition-specific inaccuracy of another credence $b(X)$ with respect to X and p is $p\mathbf{L}(X, b(X), 1) + (1 -$

---

[5]Definitions of strict propriety vary slightly between different papers - I take it that choosing this definition will not affect any of the substantive points that I make below.

[6]Pettigrew uses the term 'scoring rule' to describe a proposition-neutral local inaccuracy measure. Pettigrew defines strict propriety as a property of such scoring rules which holds if and only if, for any number $p \in [0, 1]$, the expression $p\mathbf{L}(X, b(X), 1) + (1 - p)\mathbf{L}(\neg X, b(X), 0)$ is uniquely minimised when $b(X) = p$.

$p)\mathbf{L}(X, b(X), 0))$. Since L is strictly proper, this quantity is uniquely minimised when $b(X) = p$. If we equate Pettigrew's idea of a credence expecting itself to be epistemically best with the property of a credence minimising its own expected local inaccuracy, then the observation is clearly correct.

Pettigrew's observation hints that strict propriety might be motivated by considerations of consistency similar to those that underpin Moore's paradox. Just as it seems absurd to say "it is raining but I don't believe that it is raining" so it would be odd for a rational agent's probabilistic belief function b_{pr} to assign a proposition a certain probability $b_{pr}(X) = p$, even though the agent knows that another belief function b'_{pr} has better expected accuracy with respect to b_{pr}. The belief function b_{pr} seems self-undermining, in a certain sense, since, when thinking about probabilistic agents, we usually require that agents' degrees of belief agree with certain expected values. A non-strictly-proper inaccuracy measure might lead to situations where rational probabilistic agents say odd things like 'it is probably going to rain but I expect it would be more accurate to say that it probably won't rain'. Legitimate inaccuracy measures, it might be argued, would never force agents to contradict themselves in this way.

In a similar argument, Joyce also appeals to the intuition that legitimate inaccuracy measures should not force probabilistic agents to self-undermine[7], articulating the intuition in the form of a principle he calls 'immodesty'. According to this principle, if a belief function represents a state of belief that is rational, given some state of the world, then it must uniquely minimise expected local inaccuracy with respect to itself and the domain-member that represents that state of the world. Joyce notes that a form of strict propriety is jointly implied by immodesty, together with another principle called 'minimal coherence'.

There therefore seems to be a tendency in the literature on strict propriety towards defending this principle on the grounds that it guards against self-undermining.

Defences of strict propriety that invoke self-undermining have a common problem, when used as part of accuracy arguments for probabilism. It is that the link between expectation and self-directed beliefs is only plausible in the case of probabilistic agents. A non-probabilistic agent's degree of belief $b(\neg X)$ in a negation need not be equal to $1 - b(X)$. Consequently, expected proposition-specific inaccuracies, which effectively use the quantity $1 - b(X)$ as a weight, need not have any particular significance for non-probabilistic agents.

This objection applies to both Pettigrew's hint and Joyce's argument. Pettigrew notes that his argument from self-undermining only applies to probabilistic agents, but does not make clear that the concept of expectation itself is hard to make sense of outside this context. Similarly, Joyce's principle of immodesty fails to restrict expectation to the probabilistic context, dictating that even non-

---

[7]See (Joyce, 2009, p.276-279). Note that despite the fact that Joyce's paper presents this argument in favour of strict propriety, its central accuracy argument ultimately does not incorporate it.

probabilistic belief functions must uniquely minimise their expected local inaccuracy in order to represent rational states of belief. However, non-probabilistic belief functions that fail to minimise expected inaccuracy do not necessarily represent self-undermining states of belief. Arguments for strict propriety based on self-undermining therefore seem hard to sustain. I am not aware of any other grounds on which it has been argued that inaccuracy measures should be strictly proper.

Continuity

A local inaccuracy measure is 'continuous' if it has the following property: for every proposition $X \in \Omega$, truth value $v(X) \in \{0,1\}$ and sequence of real numbers $b_1(X), \ldots, b_n(X)$ converging to $b(X) \in [0,1]$ it is the case that $\lim_{n \to \infty} \mathbf{L}(X, b_n(X), v(X)) = \mathbf{L}(X, b(X), v(X))$[8]. The requirement of continuity amounts to a stipulation that, for each proposition X, small changes in degree of belief in X lead to small changes in proposition-specific inaccuracy with respect to X.

Continuity is a problematic legitimacy condition because there are seem to be legitimate non-continuous local inaccuracy measures. One example can be constructed based on the intuition that it is a lot worse, epistemically, to hold that a false proposition is more plausible than not, even only slightly, than to be ambivalent about it. As a result, believing a false proposition to a degree slightly greater than $1/2$ should arguably make one substantially more inaccurate than if one had believed it to degree exactly $1/2$.

This intuition suggests that local inaccuracy measures should assign substantially higher scores to beliefs of degree slightly greater than $1/2$ than to than beliefs of degree exactly $1/2$. I include a graphical representation of such a proposition-specific inaccuracy measure below.

---

[8]Definition from (Landes, 2015, p.10).

Discontinuous L
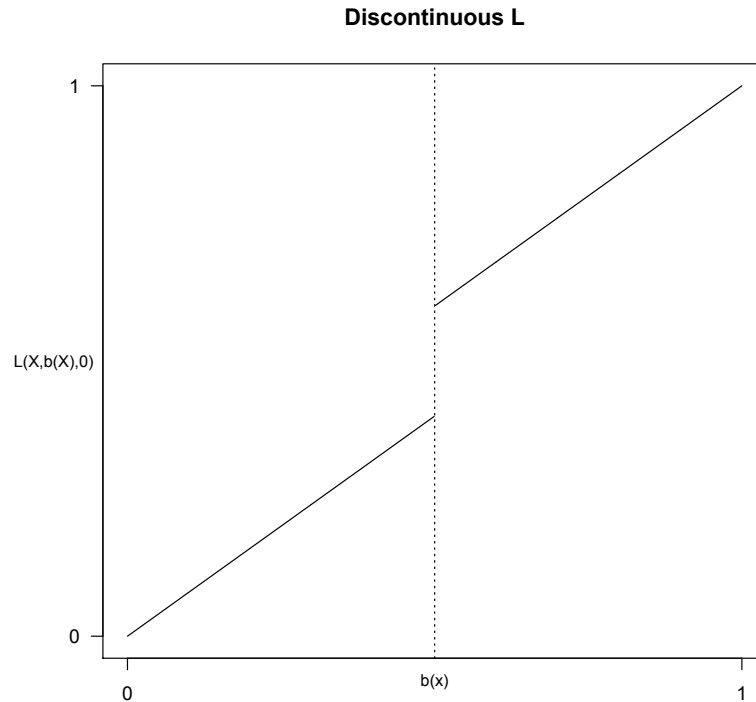
The requirement of continuity rules out this kind of local inaccuracy measure: it therefore seems to be incompatible with a plausible intuition about inaccuracy.

Conclusion

Many prominent accuracy arguments assume that legitimate global inaccuracy measures are sum-decomposable and that local inaccuracy measures are proposition-neutral, strictly proper or continuous. At least one of these assumptions is part of every currently-defended accuracy argument for probabilism that I am aware of.[9] As we have seen, Leitgeb and Pettigrew's accuracy argument assumes sum-decomposability; it also assumes continuity[10]. Though these authors do not claim to establish probabilistic necessity, the argument in Predd et al. (2009) shows that sum-decomposability, strict propriety, proposition neutrality and continuity jointly entail probabilism. Arguments for probabilistic necessity based on this demonstration need to make these assumptions as well. The ac-

---

[9]The accuracy argument in Joyce (1998) does not make any of these assumptions, but is highly controversial for other reasons: see Maher (2002) for discussion. Joyce no longer defends this argument.

[10]See (Leitgeb and Pettigrew, 2010, § 5.2.3).

curacy argument in Joyce (2009) does not assume sum-decomposability[11] but falls short of establishing that all non-probabilistic belief functions are accuracy dominated with respect to all legitimate inaccuracy measures.

The discussion above shows that all of these assumptions about the measurement of inaccuracy are problematic. Sum-decomposability has not been argued for in sufficient detail, proposition-neutrality conflicts with intuitions about evidence, strict propriety lacks a convincing rationale and continuity rules out potentially plausible inaccuracy measures.

Together with the general objections to accuracy arguments outlined above, the problems with these assumptions show that such arguments cannot yet establish probabilistic necessity.

4.2.4 Are the different arguments for probabilistic necessity mutually reinforcing?

I have argued above that, considered individually, neither betting arguments, Cox's axiomatic argument nor accuracy arguments establish that all rational states of belief can be represented by probability functions. This does not quite settle the matter of whether these arguments establish probabilistic necessity, however. It is possible that, considered collectively, the different arguments combine into a convincing argument. Views along these lines are expressed at (Gelman et al., 1995, p.13-14) and (Paris, 1994, Ch. 3). Both works argue that, while each individual argument for probabilistic necessity might not be fully convincing, the overall case is strong because the different arguments reinforce one another.

Whether the different arguments can combine into a strong argument for probabilistic necessity seems to depend on how much they have in common. The greater the extent to which the arguments are genuinely different, the more surprising it would be that they have the same conclusion, and the stronger the circumstantial support they would provide to probabilistic necessity. See Polya (1954) for a discussion of this kind of reasoning.

In this section I argue that the three kinds of argument for probabilistic necessity have many important features in common. Overall the differences between the justifications are not large enough for the arguments to combine into a compelling collective case for probabilistic necessity.

Similar assumptions

The arguments for probabilistic necessity that we have considered share many assumptions. Each assumes that rational states of belief are, in one form or

---

[11]See (Joyce, 2009, p.208).

another, classical, sharp and logically omniscient. In addition, many accuracy arguments make assumptions about aggregating quantities by adding them that are analogous to an assumption that appears in betting arguments.

Classicality

Each of the arguments for probabilistic necessity appeals to situations in which rational states of belief concern the sentences of a propositional language, and the true state of the world is described by a classically admissible valuation function.

Both of these assumptions are problematic. Propositional languages have certain limitations: for example, they cannot express reasoning about general statements like 'all cats are mammals', yet rational states of belief sometimes concern such statements. First-order languages with quantifier symbols are required to express this kind of reasoning. While there is a version of the classic Dutch-book argument for first-order languages[12], and first-order analogues of Cox's argument and accuracy arguments might be found, this does not solve the problem. Rational states of belief plausibly concern statements whose logical features even first-order languages cannot adequately describe, such as generalisations over properties.

Secondly, it is not clear that all aspects of the true state of the world can be described by a classically admissible valuation function. One problem with this assumption is that classical logic cannot easily deal with vagueness. Suppose that the sentence H expresses the proposition that a collection of fifty grains of sand forms a heap. Since all classically admissible valuation functions have the range $\{0, 1\}$, there are only two possibilities: either $v(H) = 0$ or $v(H) = 1$. Arguably, however, there are at least three possibilities as to the truth of the proposition in question: the pile either makes a heap, does not make a heap, or sort of makes a heap. See Colyvan (2004) for further arguments in favour of this and other difficulties with classicality.

The natural way to respond to concerns about classicality is to restrict the scope of the corresponding arguments. Even if esoteric phenomena like generalisations or vagueness frustrate the standard arguments for probabilistic necessity, it might be thought, these phenomena are mostly absent from the contexts where we want to evaluate peoples' partial beliefs for rationality.

Unfortunately, this pragmatic response is not available to the defender of probabilistic necessity, who must hold that every single rational state of belief can be represented by a probability function. A restricted form of probabilistic necessity—for example, one that only applied to situations that do not feature vagueness—would open the door to non-probabilistic approaches to formalising inductive reasoning in esoteric circumstances.

Consequently, probabilistic necessitarians must demonstrate that phenomena

---

[12]See (Paris and Vencovská, 2015, Ch. 5).

which appear to frustrate classicality do not actually do so, and that all epistemologically relevant aspects of belief can be described in a framework where viable arguments for probabilistic necessity can be formulated.

This task seem very difficult, and the difficulty is faced by all three kinds of argument for probabilistic necessity.

Sharp degrees of belief
According to assumptions of sharpness, every rational state of partial belief can adequately be represented by a single belief function that associates sentences with single positive real numbers. This stipulation rules out imprecise approaches to formalising uncertain reasoning, in which sets of belief functions represent states of partial belief.

The question of whether all rational states of belief are sharp is the subject of an ongoing debate in the formal epistemology literature. Elga (2010), for example, argues in favour of sharpness, whereas Bradley and Steele (2014) and Chandler (2014) dispute Elga's argument and Joyce (2010) argues that sharpness sometimes leads to irrationality.

I shall not attempt to settle this argument here. Nonetheless, it is clear that the arguments for probabilism share a contentious assumption.

Logical omniscience
In order to avoid being irrational according to probabilistic necessity, one must ensure that one's state of belief can be represented by a probability function. However, doing this can be very difficult because it requires extensive logical knowledge.

Specifically, in order to ensure that the belief function representing one's state of belief satisfies the condition $P1$, assigning the value 1 to each logical truth, one must consider every logical truth in the relevant domain. There might be very many of these.

Moreover, $P2$ requires all logical falsehoods to be assigned the value 0 and that logically equivalent sentences are assigned equal values. Even more checking is required!

Unless an agent knows all of the logical truths and falsehoods of the language that represents the objects of their belief, that is unless they are 'logically omniscient', they will surely fail to ensure that their state of belief is represented by a probabilistic belief function and thereby risk being judged irrational.

However, it can be difficult to know, in general, which of a language's sentences are logically true and false. See Parikh (1995) for a discussion of exactly how difficult it is. This difficulty provokes doubts about whether having a probabilistic state of belief is really a rationality requirement. For example, Hacking (1967)

notes that, because of the cost of computation, it is sometimes not appropriate to work out all relevant logical truths before making a decision.

A possible response to the problem of logical omniscience is to deny that rationality constraints like probabilistic necessity need to be attainable by real, non-logically-omniscient beings. Someone making this response might point out that the reach of some similar evaluative concepts extends outside the realm of possible fulfilment. For example, consider the evaluative concept of being a good digit-rememberer. Most people would have no problem saying that someone who knew the first ten billion digits of π would be a better digit-rememberer than someone who could only recall five billion, even though both feats are unattainable. Clearly it is possible to apply the concept of good remembering even in such hypothetical, unattainable, cases. Analogously, it might be argued, people with probabilistic states of belief, if they existed, would be more rational in a certain respect than non-probabilistic believers.

This analogy does not work, in my opinion, because the arguments for probabilistic necessity assert that *all* rational states of belief exhibit logical omniscience. According to the arguments, the form of rationality that representability by a probability function embodies is not instantiated by any real states of belief. In contrast, the concept of good digit-remembering applies to real rememberers as well as to hypothetical ones. In order for it to be a useful concept, there should surely be *some* real-world cases of rationality.

In order to meet the objection from logical omniscience, the defender of probabilistic necessity would therefore need to show that the rationality required to ensure representability by a probability function is an extension into the hypothetical realm of a less demanding kind of rationality that is instantiated in the real world. It is not clear that this can be done.

Sum-aggregation

Betting arguments and accuracy arguments both assume that certain quantities should be aggregated using summation rather than by other means.

As we saw in the above discussion of the packaging objection, the classic Dutch book argument assumes that the aggregate value of a package containing multiple bet-outcomes is the un-weighted sum of the values of the contents of the package. Similarly, accuracy arguments usually assume that the global inaccuracy of a state of belief is the un-weighted sum of the local inaccuracies of the degrees of belief it assigns to all propositions. Although values of bets and proposition-specific inaccuracies are different in nature, their aggregates play similar roles in the respective arguments. Betting arguments effectively claim that betting functions should not be dominated with respect to aggregate bet-value[13], whereas accuracy arguments claim that belief functions should not be dominated with respect to aggregate proposition-specific inaccuracy.

---

[13]The full Dutch book theorem, as proved in Kemeny (1955), entails that every non-probabilistic betting function is dominated by a probabilistic one.

There is thus a structural similarity between betting arguments and accuracy arguments. Admittedly, aggregating the values of bets and aggregating local inaccuracies are very different activities, and different arguments surely apply in each case. Nonetheless, someone who objects to the sum-aggregation of penalties in general might object to both arguments for the same reason.

Appeals to different kinds of reason

It is sometimes claimed that the arguments for probabilistic necessity are different because they reveal different kinds of reasons in virtue of which it is irrational to have a non-probabilistic state of belief. In this section I argue that these differences are not so clear-cut as might be thought.

The classic Dutch book argument is often taken as showing that rational agents should have probabilistic states of belief for pragmatic reasons, in contrast to the purely epistemic reasons that accuracy arguments, and perhaps also Cox's argument, purportedly unearth. For example, Joyce writes:

> ...the Dutch Book Argument ...establishes conformity to the laws of probability as a norm of prudential rationality by showing that expected utility maximizers whose partial beliefs violate these laws can be induced to behave in ways that are sure to leave them less well off than they could otherwise be. This overemphasis on the pragmatic dimension of partial beliefs tends to obscure the fact that they have properties that can be understood independently of their role in the production of action.
> (Joyce, 1998, p. 576)

However, understood properly, betting arguments are not only relevant to prudential rationality.

It is true, as Joyce says, that non-probabilistic agents can be induced to behave in ways that are sure to leave them worse off: this would be the case if the classic Dutch book scenario occurred as a real-life situation. However, situations can also arise in which probabilistic agents are sure to make themselves worse off, like the Czech book scenario from Hájek (2008). Both possibilities—actual instantiations of the Czech and Dutch book scenarios—are extremely unlikely. Thus, as far as potentially instantiated hypothetical scenarios are concerned, there does not seem to be much difference between how vulnerable probabilistic and non-probabilistic agents are to sure losses. If the classic Dutch book argument's only significance was to demonstrate such vulnerability, it would clearly be un-compelling.

However, Joyce's assertion can be interpreted differently. Perhaps he thinks that the classic Dutch book argument is only relevant to prudential irrationality because this is all that it is possible to infer from a pathological betting disposition. Recall that, in the classic Dutch book argument, an agent's choice

of a bookable betting function is taken to indicate that they have a pathological betting disposition, and this fact is then taken to show that their state of belief is irrational. The crucial question, then, is whether it is possible to infer non-prudential irrationality from a person's pathological betting disposition.

I think that this is possible. Suppose I observe my friend's behaviour at a roulette table, noticing that they tend to bet unreasonably ambitiously after they lose, while betting unreasonably cautiously immediately after winning. I take my friend's pathological betting behaviour to indicate that they are a victim of the gambler's fallacy: they hold a mistaken belief about the gambling apparatus, according to which wins are more likely than usual immediately after losses. I think that my friend is irrational—they ought to be able to work out that the roulette table does not work in this way—but the irrationality is non-prudential, as their behaviour is sensible, given their belief. My friend is simply mistaken about a matter of fact.

Similarly, it is possible to interpret the pathological betting dispositions that the classic Dutch book scenario purportedly unearths as evidence of non-prudential irrationality. Thus, betting arguments like the classic Dutch book argument can, like accuracy arguments, reveal non-prudential irrationality.

As a result, one putative difference between betting arguments and other kinds of argument for probabilism is not as stark as might have been thought.

Discussion

The different arguments for probabilistic necessity have important differences. Betting arguments appeal to hypothetical scenarios, accuracy arguments make descriptive claims about the nature of inaccuracy, while axiomatic arguments use none of these techniques. However, overall it seems that the differences between these kinds of argument are not as significant as their many similarities. All three kinds of argument make similar problematic assumptions about rational states of belief. Betting arguments and accuracy arguments depend on similarly unwarranted appeals to summation as an aggregation method. It is not clear that the different kinds of argument really appeal to different kinds of irrationality.

As a result of these similarities, the different arguments for probabilistic necessity do not combine into a compelling collective argument. It is not surprising that these arguments point in the same direction, given how similar they are: as a result, this fact does not lend their joint conclusion much additional support.

4.2.5 Conclusion: probabilistic necessity has not yet been established

I have considered three kinds of argument for probabilistic necessity, the claim that all rational states of partial belief can be represented by probability functions: betting arguments, axiomatic arguments and accuracy arguments.

We have seen that each kind of argument has specific problems that prevent it from providing a compelling standalone case for probabilistic necessity. I argued that agents without pathological betting dispositions can sometimes choose bookable betting functions in the classic dutch book scenario, that some of Cox's axioms are not compelling and that currently available accuracy arguments make implausible claims about inaccuracy and rationality.

In addition, the different arguments share many problematic features and so do not reinforce each other very much. As a result, these arguments do not combine into a strong joint argument for probabilistic necessity.

This discussion has not covered all possible arguments for probabilistic necessity. There may be other kinds of argument that establish it. Within each of the categories that I have considered, there may be alternative versions that do not suffer from the objections I have highlighted. Consequently I do not claim to have shown that probabilistic necessity is false. Nonetheless, I think it is fair to say, in light of the preceding discussion, that probabilistic necessity is epistemologically controversial.

Subjective Bayesian inductive logic therefore has a dubious commitment. This is a major problem for this research programme, especially compared to Carnapian inductive logic, which makes no comparably controversial epistemological claim.

It might be thought that commitment to probabilistic necessity is a price worth paying in order to benefit from the technical fruitfulness of being able to model degrees of plausibility as being probabilistic. This ability is indeed very useful: the success of Bayesian methods in statistics is good evidence for this. However, there is no need to commit to probabilistic necessity in order to use the probability axioms in this way. On the contrary, just as it can use other axioms to formalise inductive assumptions about, say instantial relevance, so Carnapian inductive logic can employ the probability axioms to formalise general assumptions about how inductive reasoning should proceed.

The next section addresses the other central epistemological claim of subjective Bayesian inductive logic, probabilistic sufficiency.

4.3 Probabilistic sufficiency

According to probabilistic sufficiency, all states of belief that can be represented by probability functions are rational. In this section I discuss three kinds of

argument for probabilistic sufficiency. The first two draw on topics that we have already encountered, namely the grue problem and accuracy arguments. A third kind of argument, due to Colin Howson, appeals to an analogy between the probability axioms and deductive logic. I argue that none of these arguments is successful: as a result, probabilistic sufficiency is a dubious commitment.

4.3.1   Grue-based arguments

One way to argue for probabilistic sufficiency takes advantage of the 'grue' paradox. Howson and Jeffrey put the point as follows:

> Increasing observational data certainly, provably, reinforces some hypotheses at the expense of others, but only if we let it by a suitable assignment of priors. We would like to think that an unbroken sequence of viewings of green emeralds reinforces the hypothesis that all emeralds are green. Unfortunately, it can equally be regarded as reinforcing the hypothesis that all emeralds are grue, which is inconsistent with the favoured hypothesis, unless we prevent it doing so by assigning appropriate prior weights.
> (Howson, 2000, p. 240)

> . . . if you are to update your judgemental probabilities by Bayes's theorem your prior distribution must include a programme of responses to the data, in advance. . . . Nelson Goodman's graceful illustration of this is a philosophical joke about the colours of emeralds: all that have been observed so far have been both (1) green and (2) *grue*, i.e., *green if observed before the time t, blue otherwise.* Our unhesitating choice of pattern (1) for prevision reflects a prior judgement that is dictated by experience only in the subjective sense that we are constituted so that in the light of that experience we are incapable of taking the "grue" hypothesis seriously.
> (Jeffrey, 1988, p. 8, emphasis original)

The following train of reasoning seems to lie behind these quotations. First, the conditional probability method for taking evidence into account is assumed. It is then noted that whenever a given probability function seems to represent the view that a certain piece of evidence supports a certain hypothesis, there is an alternative, grue-like interpretation of its domain according to which the same probability function represents the view that that evidence undermines the hypothesis.

Something along these lines seems to be true. For example, suppose that we think that the evidence that a large sample of n emeralds examined before the crucial time t and found to be green supports the hypothesis that the first emerald examined after t will be found to be green. We decide to represent the evidence with a sentence $Ga_1 \wedge \ldots \wedge Ga_n = E \in SL_{pred}$ and the hypothesis with the sentence $Ga_{n+1} = H \in SL_{pred}$. We depict our view that the sample in

question supports the hypothesis by representing our beliefs with a probability function $pr^\star : SL_{pred} \to [0,1]$ such that $pr^\star(H \mid E) > pr^\star(H)$. To our disappointment, however, we find that, had we chosen an interpretation of SL_{pred} according to which G represented grueness rather than greenness, H would have represented the hypothesis that the first emerald examined after t is blue, not green. Since it would still have been the case that $pr^\star(H \mid E) > pr^\star(H)$, under this interpretation the probability function $pr^\star$ would have represented the view that our sample undermined the hypothesis, rather than the view that it supports it.

Analogously, the argument goes, any probability function that represents a particular strategy according to a familiar interpretation of its domain represents a completely different strategy strategy according to a grue-like interpretation. The grue problem therefore illustrates Hume's point that all strategies for making predictions on the basis of evidence are equally unsound. Since every probability function represents a strategy for taking evidence into account, the fact that all such strategies are, in a sense, equal suggests that all probability functions should be considered equal as well.

As we saw in the previous chapter, the grue problem shows that which inductive assumption many inductive logical axioms represent depends crucially on how the domains of the measure functions they govern are interpreted. According to proponents of the grue-based argument, the probability axioms are not interpretation-sensitive in this way, but rather represent the same, plausible, assumptions under all interpretations of their measure functions' domains. This fact, they conclude, shows that only the probability axioms represent genuine rationality constraints.

The key claim of this argument is that axioms that formalise genuine rationality principles are not interpretation-sensitive. If an axiom formalises a rationality constraint under one interpretation, the argument claims, it should do so under *every* interpretation. I shall now argue that this claim is false: on the contrary, there seem to be genuine, interpretation sensitive rationality principles.

It is widely accepted that responding appropriately to one's evidence is important for maintaining one's rationality. In order to respond appropriately to one's evidence one must distinguish propositions that are evidentially supported from those that are not, and treat these two categories of propositions differently. Formal rationality principles encoding what constitutes a reasonable response to evidence—call them 'evidence principles'—must consequently impose a similar distinction between formal objects that represent evidentially-supported propositions and formal objects that represent other propositions. However, any principle whose content depends on which propositions particular formal objects represent is interpretation-sensitive. Thus the principle that one should respond appropriately to one's evidence seems to be interpretation sensitive.

To illustrate this point, consider the Principal Principle, an evidence principle that attempts to represent what constitutes appropriate reasoning when one

has evidence about chance propositions. The Principal principle applies to elements of rich domains D_{rich} with the form $\langle Ch(H) = p \rangle$, which are intended to represent propositions about physical chances, namely that the proposition represented by H has certain physical chance p. According to the Principal principle, if $b_{pr} : D_{rich} \to [0,1]$ represents a rational state of belief, then $b_{pr}(H \mid \langle Ch(H) = p \rangle) = p$. Clearly the Principal principle has very different content depending on how the members of D_{rich} are interpreted. If elements like $\langle Ch(H) = p \rangle$ are interpreted as intended, then the Principal principle arguably represents a rationality requirement. However, if, for example, the element $\langle Ch(H) = p \rangle$ is interpreted as representing the proposition that H has physical chance $p/2$, then unless $p = 0$, the principal principle represents a clearly irrational attitude. Thus the Principal principle is interpretation-sensitive.

A defender of interpretation-neutrality might argue that this is not a problem, as it merely shows that the principal principle does not formalise a genuine rationality constraint. In any case, they might continue, the principal principle is highly controversial - what could be so bad about taking the perfectly defensible position that it is not a rationality constraint?

This response avoids the main issue, however. Interpretation-neutrality does not just rule out the principal principle, but seemingly any principle that distinguishes evidentially supported propositions from evidentially undermined ones. As a result the pertinent aspect of the example is not the principal principle's substantive content, but the fact that it is interpretation-sensitive. While the former is indeed controversial, it is not usually argued that the principal priciple's interpretation-sensitivity precludes the possibility of its being a genuine rationality constraint.

To sum up, Grue-based arguments for probabilistic sufficiency depend on the assumption that all genuine rationality principles are interpretation-neutral. However, there seem to be at least some formal principles that formalise how one should rationally take evidence into account, and it seems likely that any such principle would be interpretation-sensitive, I conclude that Grue-based arguments for probabilistic sufficiency are unsucessful.

4.3.2   Dominance-avoidance

Another way to argue for probabilistic sufficiency draws on accuracy arguments for probabilism, which we encountered in section 4.2.3 above. This kind of argument for probabilistic sufficiency is particularly noteworthy as one was advanced by De Finetti in De Finetti (1990).

Many accuracy arguments for probabilism not only show that only probability functions avoid accuracy-domination, but also that every probability function is non-dominated. In other words, if one grants accuracy arguments' preliminary assumptions and claims about legitimate inaccuracy measures, then one must

accept that representability by a probability function is sufficient for avoiding accuracy-domination. Predd et al. (2009) contains a result along these lines.

This kind of result can be turned into an argument for probabilistic sufficiency by assuming the epistemological principle, which I shall call *Dom*, according to which every belief function that is not accuracy-dominated represents a rational state of belief.

We saw above that many of the assumptions underlying accuracy arguments for probabilistic necessity are suspect: these objections also apply to accuracy arguments for probabilistic sufficiency.

Moreover, accuracy arguments for probabilistic sufficiency have an additional problem: just like interpretation-neutrality, *Dom* frustrates plausible evidence-based rationality principles.

Suppose that, according to a certain body of evidence, a certain state of the world is possible but very unlikely; for example, suppose that, given the evidence, it is extremely implausible that unicorns exist. Suppose that, contrary to the evidence, agent A assigns a very high degree of belief to the proposition that unicorns exist. Given their extreme defiance of the evidence, we might be tempted to conclude that agent A is irrational. This would remain the case even if it turned out that that there were, in fact, unicorns - this would be a case of mere epistemic luck rather than sound judgement, and so ought not affect our judgement that A is irrational.

However, if the unicorny state of affairs that they believe in were true, then we would have to conclude that, at least as far as this proposition is concerned, agent A is very accurate. Since the existence of unicorns cannot be ruled out altogether, a belief function representing agent A's state of belief will never be accuracy-dominated despite defying all the available evidence. Therefore, if we accept *Dom*, we seem forced to conclude that agent A's state of belief is rational.

Due to the problems that accuracy arguments for probabilistic sufficiency share with accuracy arguments for probabilistic necessity, together with the implausibility of *Dom*, I conclude that such arguments cannot establish probabilistic sufficiency.

4.3.3 Howson's logical argument

In (Howson, 2000, Ch. 7) and several related articles—Howson (2003), Howson (2012), Howson (1997a), Howson (2001), Howson (1997b)—Howson argues that the probability axioms are "complete". Unlike stronger collections of axioms, he claims, the probability axioms can underpin a discipline with properties that Howson takes to characterise "genuine" logic. Insofar as axioms that represent rationality principles must be genuinely logical in this sense, Howson's argument can be used to support probabilistic sufficiency.

In this section I investigate Howson's proposed criteria of genuine logicality. I argue that, depending on how we interpret what Howson writes, genuine logicality either fails to distinguish the probability axioms from stronger collections of axioms or else does not constitute a compelling standard of rationality. In either case, Howson's argument fails to establish probabilistic sufficiency.

Howson's argument

Howson's argument begins with the following proposed criteria of genuine logicality. Satisfying all of these criteria, he claims, is necessary and sufficient for any discipline to be a genuine logic:

$GL(a)$ It involves statements and relations between them.

$GL(b)$ It adjudicates some mode of non-domain-specific reasoning.

$GL(c)$ It is 'about consistency'. More specifically, it should incorporate a semantic notion of consistency which can be shown to be extensionally equivalent to one characterizable purely syntactically; this equivalence is the content of what are called soundness and completeness theorems for the corresponding system.

(Howson, 2000, p. 127, 'GL's added)

The first criterion, $GL(a)$, is satisfied by all putative logics that we will investigate and so does not require further investigation. $GL(b)$ requires genuine logics to govern modes of reasoning that are not 'domain-specific'. Howson does not fully explain what makes a mode of reasoning domain-specific, though he does mention that Bayesian probability satisfies $GL(b)$ because "any proposition whatsoever can be in the domain of a probability function" (Howson, 2000, p. 127). This suggests the following fleshing out of Howson's criterion: genuine logics must be capable in principle of representing reasoning about arbitrary propositions.

The final criterion, $GL(c)$, stipulates that genuine logics must involve extensionally equivalent syntactic and semantic notions, and that these must be notions of consistency. Howson claims to have found a logic of uncertain reasoning that satisfies all of these criteria.

Howson takes the probability axioms to provide this logic's syntactic notion of consistency: a measure function is syntactically consistent if and only if it is a probability function.

Howson obtains a semantic notion of consistency by interpreting measure functions as possible actions in a betting scenario. A measure function is semantically consistent, according to this notion, if and only if it represents a 'fair' bet. Howson proposes semantic conditions that, he claims, characterise fairness.

Informally, for a bet on a proposition A to be fair in Howson's sense means that, given some possible state of knowledge, it is equally advantageous to take either side of the bet. More formally, Howson claims at (Howson, 2003, p. 159) that it makes sense to equate bets with 'betting quotients' $p(A) \in [0, 1]$, representing the proportions into which a disinterested but principled third party might think it would be appropriate to split some wagered money if the truth of A has not yet been revealed. $p(A)$ represents the proportion of the money that the arbiter gives to the person who stands to gain the whole lot if A is revealed to be true, and $1 - p(A)$ the proportion that is given to the person who takes the other side of the bet.

Such betting quotients are fair, Howson claims, just when they satisfy the following conditions:

> H1 If p is a fair betting quotient then $0 \leq p \leq 1$
>
> H2 If A is a logical truth then $P(A)$, the fair betting quotient for A, is 1, and... if A is a logical falsehood $P(A) = 0$.
>
> H3 Fair bets must obviously be invariant under change of sign of the stake: if the bet on A has no advantage over the bet against A, the converse must hold too.
>
> H4 If a fair bet on some proposition A, or a sum of fair bets on some subset of propositions, determines a bet on some proposition B with stake S and betting quotient q, then... p [is] the fair betting quotient on B.
>
> (Howson, 2000, p. 129, 'H's and square parentheses added)

Howson proves that these two notions of consistency are equivalent in this sense: all measure functions that represent fair bets are probability functions, and all probability functions represent fair bets. Consequently, probabilism is sound and complete, in Howson's sense, with respect to his notion of fairness, as spelled out by $H1 - 4$. On this basis Howson concludes that the probability axioms underpin a genuine logic.

Furthermore, Howson argues, as a consequence of the probability axioms' completeness relative to his semantics, super-probabilistic collections of axioms cannot be part of genuine logics of uncertain reasoning. This can be seen from the following quotations:

> According to the Bayesian theory the probability axioms express a complete *logic* of credibility judgements.
> (Howson, 2000, p. 178, emphasis original)

> [the logical view of the principles of subjective probability] also tells us what the inductive premise will look like: it will be a probability assignment that is not deducible from the probability axioms.
> (Howson, 2000, p.134)

> Under the aspect of logic the probability axioms are as they stand
> complete: they are, with a qualification we shall discuss in Chapter 9,
> a complete set of constraints for consistency. Hence any extension of
> them–as in principles for determining 'objective' prior probabilities–
> goes beyond pure logic.
> (Howson, 2000, p.132)

To summarise, then, Howson claims that the probability axioms are the strongest
genuinely logical axioms for measure functions. Unlike stronger collections of
axioms, he claims, the probability axioms can provide the basis of a discipline
that satisfies his requirements of genuine logicality $GL(a) - (c)$.

Response: an alternative logic of uncertain reasoning

I shall now object to this method of arguing for probabilistic sufficiency. The
probability axioms, I claim, either are not the strongest inductive logical axioms
that can underpin a genuinely logical discipline, or else axioms that are not
genuinely logical can nonetheless express rationality principles.

In order to make this case I will attempt to replicate Howson's procedure for a
stronger collection of inductive logical axioms. Following Howson, this means
that I shall need to find syntactic and semantic notions of consistency and show
that they are extensionally equivalent, and then show that the the discipline
drawing on these notions, which I call the 'alternative logic', satisfies conditions
$GL(a)$ and $GL(b)$. I will claim that, depending on how we interpret Howson's
criterion $GL(b)$, the alternative logic is either genuinely logical according to
his criteria, or else shows that the criteria can be contravened by rationality
principles. In either case, it will be clear that Howson's argument does not
establish probabilistic sufficiency.

Syntactic notion of consistency

The alternative logic's syntactic notion of consistency is given by the collection
of axioms consisting of the probability axioms and also the axiom of constant
exchangeability **Ex**.[14] This collection of axioms is super-probabilistic because
it is strictly more restrictive than the probability axioms.

Semantic notion of consistency

The alternative logic's semantic notion of consistency interprets measure func-
tions as determining betting quotients relating to the game 'Spinner'. In this
game an arrow is spun around on a pin with the end pointing at some coloured
circle-sectors. The result of a round of the game is determined by the final rest-
ing point of the end of the arrow. If the apparatus is set up correctly then the
results of previous rounds are uninformative as to the result of the next round.

---

[14]It is assumed in this section that the domain of a probability function is the set of sentences
SL_{pred} of a unary predicate language.

Just as in Howson's betting scenario, a disinterested third party is asked to divide a stake that has been bet on a certain proposition A. In this case, however, the proposition A specifies a future sequence of outcomes of the spinner.

Suppose that $H1 - H4$ are, as Howson claims, appropriate semantic conditions for betting in general; as a result they would also apply to the specific case of bets on spin outcomes. Furthermore, I claim that the following semantic principle is also a fairness condition for Spinner moves:

H5 If A is a proposition specifying a certain sequence of spin outcomes, and B is a proposition specifying a sequence that is a permutation of of the sequence specified by A, then a fair betting quotient for A should also be a fair betting quotient for B.

This condition is necessary for fair division of stakes in the Spinner scenario. If it could be sensible to make a different division of stakes for A than for B, then the spins would not be unpredictable in the way required for fair play. $H5$ is therefore surely at least as compelling as $H1 - H4$.

We can now formulate the alternative logic's semantic notion of consistency. A measure function is semantically consistent, according to the alternative logic, if and only if it represents betting quotients for spin outcomes that satisfy the principles $H1 - H5$.

Soundness and completeness
The syntactic notion of consistency specified by the probability axioms, together with **Ex** is sound and complete with respect to the semantic conditions characterising fair bets on spin outcomes in exactly the same way that the probability axioms are sound and complete with respect to the semantic conditions that characterise Howson's betting game. By Howson's argument, $H1$-$H4$ guarantee soundness and completeness with respect to probabilism, while $H5$ merely restates the definition of **Ex** according to the interpretation of sentences as spinner sequences specified above.

The alternative logic therefore satisfies Howson's criterion $GL(c)$. It also clearly involves statements and relations between them: Howson's condition $GL(a)$ is therefore satisfied.

Domain-specificity
It remains to be seen whether the alternative logic satisfies Howson's criterion $GL(b)$, that is, whether it "adjudicates some mode of non-domain-specific reasoning". This task is difficult because, as mentioned above, Howson is not entirely clear as to what $GL(b)$ entails.

Drawing on Howson's apparent claim that his probability-based logic is non-domain-specific because "any proposition whatsoever can be in the domain of a probability function", we might be tempted to conclude that the alternative logic

is non-domain-specific as required. After all, just as a member of a probability function's domain can be interpreted as representing any proposition, the same is true of a member of the domain of a probability function satisfying **Ex**. If this were all there is to $GL(b)$, then the alternative logic would seem to satisfy this condition, and would therefore be a genuine logic, according to Howson's criteria.

However, although the axioms that constitute the alternative logic's syntax might just as non-domain-specific as the probability axioms, the alternative logic's semantics of spinner bets seem to be limited in a way that Howson's semantics of betting in general are not. Whereas Howson's conditions $H1-4$ are arguably requirements of fair betting on any proposition whatsoever, $H5$ is only a fairness requirement for bets on propositions that specify outcomes of the spinner, or else are analogous to such propositions.

This difference might make one think that, even though the alternative logic can represent reasoning about any proposition, it does not truly adjudicate a mode of non-domain-specific reasoning because one of its semantic conditions only applies within a certain domain of propositions. As a result, one might continue, the alternative logic fails to satisfy Howson's criterion $GL(b)$.

This reading of $GL(b)$ suggests that any discipline whose betting-based semantic conditions do not apply to bets on arbitrary propositions cannot be a genuine logic. Whether or not it is a requirement of genuine logicality, satisfying this form of $GL(b)$ is surely not a requirement for axioms representing genuine rationality principles. Some axioms seem to represent genuine rationality principles, yet fail this test.

Consider again the principal principle, which was introduced above in section 4.3.1. As mentioned there, it arguably formalises a rationality principle, so long as terms with the form $\langle Ch(H) = p \rangle$ are interpreted as intended, as representing the right kind of chance proposition. However, as we have seen, the principal principle is interpretation-sensitive: as a result it seems that any logic based on the principal principle will be domain-specific and therefore fail to satisfy $GL(b)$. Even if the principal principle is not accepted as a genuine rationality requirement, we saw above that it is difficult to capture rational evidence-based reasoning without interpretation sensitivity. It therefore seems that, on the second reading, $GL(b)$ is not satisfied by all genuine rationality principles.

To sum up, if we interpret Howson's condition $GL(b)$ in the first way, then it is too liberal. On this reading the alternative logic is genuine according to Howson's criteria, and the probability axioms are not, as he claimed, the strongest axioms that underpin a genuine logic. On the other hand, if we interpret $GL(b)$ in the second way, then it is too restrictive. On this reading, there are genuine rationality principles that fail to underpin genuine logics. On both readings, Howson's arguments cannot be used to establish probabilistic sufficiency.

4.3.4 Probabilistic sufficiency has not been established

Neither grue-based arguments, accuracy arguments nor Howson's logic-based argument satisfactorily established probabilistic sufficiency. Grue based arguments depend on the dubious assumption that axioms representing rationality principles are interpretation-neutral. Accuracy arguments, among other problems, require all non-accuracy-dominated belief functions to represent rational states of belief. This requirement is hard to square with the fact that irrational states of belief can be accurate due to luck. Finally, Howson's logic-based argument either fails to establish that super-probabilistic collections of axioms cannot underpin genuine logics, or else prevents some axioms that represent rationality principles from doing so.

As each of these arguments for probabilistic sufficiency is unsuccessful, and these are the only such arguments that I am aware of, I conclude that a compelling case for probabilistic sufficiency has not yet been made.

4.4 Which kind of inductive logic is best?

This section directly compares subjective Bayesian and Carnapian inductive logic. I argue that Carnapian inductive logic is preferable to subjective Bayesian inductive logic because its philosophical presuppositions are less controversial, and because it is more versatile.

4.4.1 Philosophical comparison

Subjective Bayesian inductive logic is committed to probabilistic sufficiency and probabilistic necessity: we have seen above that both of these epistemological claims are dubious. It is also committed to the view that rational agents take evidence into account using the conditional probability method for representing evidence, though I will not discuss the epistemological merits of this commitment.

Carnapian inductive logic makes none of these epistemological claims. It is uncontroversial that Carnapian inductive logic was not committed to probabilistic sufficiency - all of the systems of inductive logic that Carnap proposed went beyond the probability axioms. Probabilistic necessity goes against Carnapian inductive logic's commitment to tolerance, as it entails rejecting some systems of inductive logic on philosophical rather than pragmatic grounds. Finally, Carnapian inductive logic was not committed to using the conditional probability method in all cases. This can be seen from passages like (Carnap, 1971a, p.120), where Carnap argues that evidence should sometimes be represented informally in the choice whether to employ axioms such as exchangeability, and Carnap

(1967b) where Carnap entertains an alternative formal method for representing uncertain evidence.

The main commitments that Carnapian inductive logic did make—to explication, systems of inductive logic and tolerance—are philosophically uncontentious in comparison to those of subjective Bayesian inductive logic.

Although some authors, notably Strawson (1963), have objected to Carnap's method of explication on general grounds, these objections can be overcome, as shown, for example, by Maher (2007). Recently several authors, including Lutz (2012), Leitgeb (2013) and Olsson (2014), have proposed explicitly explicative projects, showing that this method is widely regarded as viable.

The suitability of systems of inductive logic as explicata for inductive assumptions is questionable in principle but far more defensible than probabilistic necessity and sufficiency. In any case, subjective Bayesian inductive logic effectively shares this commitment. It too proposes to represent certain inductive assumptions—those of rational agents—with certain systems of inductive logic, namely systems $(D, \{pr\})$ where the members of D represent possible objects of partial beliefs and pr represents the credences of a rational agent.

Finally, Carnapian inductive logic's commitment to tolerance, according to which systems of inductive should be evaluated according to their practical rather than philosophical merits, is surely far less controversial than the commitments of subjective Bayesian inductive logic.

It seems, then, that subjective Bayesian inductive logic is more philosophically contentious than Carnapian inductive logic.

4.4.2   Methodological comparison

Subjective Bayesian inductive logic is undeniably simpler than Carnapian inductive logic. Rather than addressing systems of inductive logic in general, it restricts its attention to the far smaller class of probabilistic systems of inductive logic, arguing that other systems fall outside the purview of pure logic. Rather than dealing with the problem of formalising inductive assumptions in general, it addresses only the task of formalising rational partial beliefs. Rather than allowing any old epistemological theory to be represented, subjective Bayesian inductive logic restricts its attention to subjective Bayesian epistemology.

If it were never practically useful for inductive logic to depart from any of these prescriptions, then subjective Bayesian inductive logic's comparative simplicity might make it more attractive than Carnapian inductive logic, in spite of any qualms about its philosophical commitments. However, this is not the case. On the contrary, the practical benefit of subjective Bayesian inductive logic's comparative simplicity is outweighed by other practical advantages of Carnapian inductive logic.

Since it countenances a far wider range of inductive assumptions, Carnapian inductive logic can be used for many tasks that subjective Bayesian inductive logic cannot be used for. For example, Carnapian inductive logic can formalise the inductive assumptions of agents with non-probabilistic states of belief. Whether or not such agents are strictly rational, they are thought to be common—see, for example, Gigerenzer and Goldstein (1996) for a discussion of this point—and as such their inductive assumptions surely merit formalisation. Similarly, Carnapian inductive logic can also address relationships between inductive assumptions that do not enter the picture within subjective Bayesian inductive logic. For example, it was shown in Gaifman (1971) that the axiom of constant exchangeability entails the axiom of non-negative instantial relevance, showing an interesting relationship between the assumptions that these two axioms formalise. Such insights fall outside the purview of subjective Bayesian inductive logic as both axioms are stronger than probabilism.

As well as being able to formalise a wider range of inductive assumptions, Carnapian inductive logic gains a practical advantage from the fact that it is not restricted in how it can express evidence. Whereas, within subjective Bayesian inductive logic, all evidence must be represented in the domain of a probability function, Carnapian inductive logic allows evidence to be represented using an alternative formal method, as in Carnap (1967b), or alternatively allows evidence to be represented informally in the choice of a system of inductive logic. For example, if the available evidence suggests that the order of experiments is inductively irrelevant, this evidence can be represented informally through the choice of a system of inductive logic that employs the axiom of constant exchangeability. We shall see in chapter 6 that this flexibility allows Carnapian inductive logic to capture the practices of working statisticians more effectively than would be practicable within subjective Bayesian inductive logic.

These examples seem to indicate that Carnapian inductive logic makes a better trade-off between generality and simplicity than subjective Bayesian inductive logic. As a result, Carnapian inductive logic seems to be preferable to subjective Bayesian inductive logic from a practical point of view.

4.4.3 Conclusion

Subjective Bayesian inductive logic makes commitments that are more philosophically contentious than those of Carnapian inductive logic, cannot formalise as great a range of inductive assumptions, and cannot represent evidence as flexibly as Carnapian inductive logic. On the other hand, subjective Bayesian inductive logic is simpler in a certain sense.

While a trenchant proponent of subjective Bayesian inductive logic might argue that the benefit of their chosen discipline's simplicity is so great that it offsets each of the points in favour of Carnapian inductive logic, I think that this stance would be unreasonable. Simplicity is not such an important feature

of an approach to inductive logic that it trumps all other philosophical and methodological factors. I therefore conclude that Carnapian inductive logic is preferable to subjective Bayesian inductive logic.

Chapter 5

Objective Bayesianism

5.1 Introduction

Objective Bayesian epistemology, or objective Bayesianism, asserts that all rational states of partial belief can be represented by probability functions, but denies that all states of belief that can be so represented are rational. In this chapter I explore two ways of connecting objective Bayesian epistemology with Carnapian inductive logic. I focus on the form of objective Bayesian epistemology presented in Williamson (2010).

The first way uses objective Bayesian epistemology to construct a distinct objective Bayesian form of inductive logic. This approach is analogous to subjective Bayesian inductive logic which, as we saw in the previous chapter, uses subjective Bayesian epistemology to construct subjective Bayesian inductive logic.

I compare objective Bayesian inductive logic with Carnapian inductive logic, arguing that Carnapian inductive logic is philosophically preferable, while both forms of inductive logic have distinct methodological advantages.

The second, 'assimilative', approach to connecting objective Bayesian epistemology and Carnapian inductive logic treats the rational norms that characterise the former as inductive assumptions to be formalised.

I describe in general terms how this can be achieved, showing how a system of inductive logic might formalise objective Bayesian epistemology, and argue that this would be a good idea.

I conclude by comparing the two approaches. While objective Bayesian inductive logic and the assimilative approach are very different, I claim that they should not be thought of as opposed to each other and might even be complementary.

5.2   Objective Bayesian epistemology

Many of the most famous proponents of objective Bayesian epistemology have not been professional philosophers, but rather philosophically minded scientists, such as Laplace, Leibniz, Jeffreys, Jaynes and Berger. Disciplinary boundaries notwithstanding, I shall include these authors in the category 'objective Bayesian epistemologists'.

Objective Bayesian epistemologists are united by the belief that a rational state of partial belief should be representable by a probability function, and also satisfy some additional constraints. Objective Bayesian epistemology therefore shares with subjective Bayesian epistemology a commitment to probabilistic necessity, but differs from it by rejecting probabilistic sufficiency. All objective Bayesian epistemologists think that some probabilistic states of belief are irrational. Some proponents of objective Bayesian epistemology support conditionalisation, whereas others reject it in favour of alternative ways of describing how rational agents should react to evidence.

I shall focus on the form of objective Bayesian epistemology defended by Jon Williamson in Williamson (2010). It would be difficult to do justice to all the authors mentioned above, and Williamson's version of objective Bayesian epistemology shares many features with the others. In addition, Williamson's position is distinguished by being particularly distinct from subjective Bayesianism. First, Williamson, like many other objective Bayesian epistemologists, supports a form of the equivocation norm, according to which rational agents ought to distribute their degrees of belief as evenly as possible. Second, and perhaps more unusually, Williamson's position also differs from subjective Bayesianism as to how rational agents ought to react to evidence.

5.2.1   Norms of rational belief

Williamson's form of objective Bayesian epistemology position imposes two norms on rational states of partial belief in addition to the probability norm. These are called calibration and equivocation. These norms are presented in Landes and Williamson (2013) as follows:

> **Probability** The strengths of an agent's beliefs should be representable by a probability function.
>
> **Calibration** The strengths of an agent's beliefs should satisfy constraints imposed by her evidence $\mathcal{E}$.[1]

---

[1]Note that, here and in what follows, I use the caligraphic symbol $\mathcal{E}$ to represent an agent's evidence, whereas the presentation in Landes and Williamson (2013) use the symbol E. This is to distinguish the symbol representing an agent's evidence from symbols representing propositions or sentences.

> **Equivocation** The agent should not adopt beliefs that are more extreme than is demanded by her evidence $\mathcal{E}$.

(Landes and Williamson, 2013, p. 1)

The probability norm forces rational states of partial belief to be representable by probability functions. It is effectively equivalent to probabilistic necessity.

The calibration norm dictates how a rational agent should take evidence into account. Every possible state of evidence $\mathcal{E}$ is taken to determine a corresponding set $\mathbb{E}_{\mathcal{E}} \subseteq \mathbb{P}$ of probability functions which are calibrated with that evidence. The calibration norm stipulates that any rational state of partial belief that takes into account $\mathcal{E}$ should be representable by a probability function in $\mathbb{E}_{\mathcal{E}}$. For the sake of readability I sometimes omit the subscript $_{\mathcal{E}}$ below when there is no need to indicate which state of evidence corresponds to a particular set of calibrated probability functions.

The problem of working out which set $\mathbb{E}_{\mathcal{E}}$ contains the probability functions that are calibrated with a given state of evidence $\mathcal{E}$ is somewhat open-ended. This is perhaps fitting, as there are many ways in which evidence can constrain the strengths of one's beliefs. Landes and Williamson present one form of constraint—information about physical probabilities—as follows:

> In particular, if the evidence determines just that physical probability (aka chance) $P^{\star}$ is in some set $\mathbb{P}^{\star}$ of probability functions... then [the probability function representing the strengths of an agent's beliefs] should be calibrated to physical probability insofar as it should lie in the convex hull $\mathbb{E} = \langle \mathbb{P}^{\star} \rangle$ of the set $\mathbb{P}^{\star}$.
> (Landes and Williamson, 2013, p.1, parentheses original)

In addition, objective Bayesian epistemologists envisage other kinds of evidential constraint. One's evidence might compel one to consider only probability functions which assign certain expected values to particular random variables. (Jaynes, 1995, Ch. 11) explains how this kind of evidence can help to determine a set of calibrated probability functions. At (Williamson, 2010, p.28) Williamson considers 'structural' constraints corresponding to evidence about things other than physical probabilities.

The equivocation norm indicates which states of partial belief among those permitted by their evidence a rational agent ought to adopt. Given a set of calibrated probability functions $\mathbb{E}_{\mathcal{E}}$ and a desired level of equivocation, the equivocation norm identifies a set $\Downarrow \mathbb{E}_{\mathcal{E}}$ of sufficiently equivocal probability functions. It stipulates that rational states of partial belief taking into account $\mathcal{E}$ should be representable by probability functions in $\Downarrow \mathbb{E}_{\mathcal{E}}$.

The precise content of the equivocation norm depends on how exactly degrees of equivocation are measured and on what level of equivocation is deemed to be sufficient. In the case of a domain Ω that is a state space generated by a finite number of atomic states, Williamson proposes that the degree of equivocation of a given probability function p should be measured by its Shannon

entropy $H(p) = -\sum_{\omega \in \Omega} p(\omega) \log p(\omega)$. Analogues of this quantity can be found for domains that are sets of sentences of propositional or unary predicate languages: see Landes and Williamson (2013) for details. The level of equivocation required for membership of $\Downarrow \mathbb{E}_{\mathcal{E}}$ is typically the maximum level possible,though Williamson allows that in some cases it might be preferable to set a very high but sub-maximal level.

I will not dwell on the issues of how to measure equivocation and what is the best level. I shall assume that a set of sufficiently equivocal calibrated probability functions can be defined uncontroversially in any circumstance.

5.2.2   Are the norms justified?

Objective Bayesian epistemology advocates the probability norm, which is tantamount to probabilistic necessity; it is therefore just as vulnerable as subjective Bayesian epistemology to the epistemological objections laid out in section 4.2 above.

Since I take it that the arguments in section 4.2 already establish that the probability norm is dubious, I will not direct very much attention towards the question of whether the calibration and equivocation norms are justified. Landes and Williamson (2013) discusses these issues at length.

5.2.3   Methods of representing evidence

The calibration norm introduces a novel 'calibration-based' way of representing the effect of taking evidence into account. According to the calibration norm, every state of evidence $\mathcal{E}$ determines a set of calibrated probability functions $\mathbb{E}_{\mathcal{E}}$. Taking evidence into account is represented by adopting a probability function in $\mathbb{E}_{\mathcal{E}}$.

This calibration-based method of representing taking evidence into account differs from the conditional probability method that subjective Bayesian epistemology employs. Like the calibration-based method, the conditional probability method begins by stipulating that rational states of belief should be represented by probability functions. However, unlike the calibration-based method, the conditional probability method represents every state of evidence $\mathcal{E}$ by a specific member $E_{\mathcal{E}} \in D$ of the relevant domain. The effect of taking $\mathcal{E}$ into account is then represented by the difference between $m_{initial}(\bullet)$ and $m_{new}(\bullet) = m_{initial}(\bullet \mid E_{\mathcal{E}})$.

Mirroring the presentation of the calibration-based method, one could formulate the conditional probability method as a norm according to which, if one's evidence is represented by $E_{\mathcal{E}}$, then one's state of belief should be represented by a belief function in the set $\mathbb{P}_{E_{\mathcal{E}}} = \{m : m \in \mathbb{P} \text{ and } m(E_{\mathcal{E}}) = 1\}$ of probability functions that assign the value 1 to $E_{\mathcal{E}}$.

The main novelty of the calibration-based method is that, except insofar as it must determine a set $\mathbb{E}_\mathcal{E}$, evidence itself need not be represented formally. In particular, states of evidence need not be assigned degrees of belief.

In contrast, the conditional probability method of representing evidence requires that all possible states of evidence must be represented formally as domain-members and therefore be assigned degrees of belief.

In addition, the calibration-based method allows a sharp distinction between considerations to do with evidence and other considerations. In Williamson's objective Bayesian epistemology, the calibration norm encompasses all evidence-related considerations, whereas the probability and equivocation norms do not have to do with evidence. The conditional probability method does not make this kind of distinction.

Agreement and disagreement between the two methods

In certain cases the two ways of representing evidence agree. For example, suppose that the only effect of taking into account the state of evidence $\mathcal{E}$ is to recognise the truth of some proposition represented by the domain-member X.

In this case, the calibrated set $\mathbb{E}_\mathcal{E}$ should plausibly contain all and only those probability functions that give probability 1 to X. This is the same as the set $\mathbb{P}_X$ recommended by the conditional probability method.

However, the two methods do not always agree, as the effect of taking evidence into account is not always to recognise the truth of propositions corresponding to domain-members. In particular, evidence can impart information about physical probabilities. Suppose, for example, that the state of evidence $\mathcal{E}_{ch}$ dictates only that the physical probability of the proposition represented by domain member A is between $1/2$ and $2/3$. The calibration-based method must represent $\mathcal{E}_{ch}$ by making sure that $\mathbb{E}_{\mathcal{E}_{ch}}$ contains only probability functions that assign A values between $1/2$ and $2/3$. On the other hand, the same effect need not occur under the conditional probability method. Even if the proposition that $\mathcal{E}_{ch}$ is taken into account corresponds to a domain-member—call it E_{ch}—the set $\mathbb{P}_{E_{ch}}$ can in principle contain probability functions that assign A values outside the interval $[1/2, 2/3]$. Unless an axiom along the lines of the principal principle is imposed, the calibration-based method and the domain-based method will therefore disagree in this case.

(Williamson, 2011, § 4) explains, in general, when the two methods agree and when they do not.

Advantages and disadvantages of each method

Not having to represent states of evidence in measure functions' domains confers several advantages on the calibration-based method for representing taking

evidence into account.

From a philosophical point of view, the calibration-based method is, in principle, liberated from holding that evidence always takes the form of a collection of propositions. Objective Bayesian epistemologists may therefore potentially allow that certain objects may constitute evidence. This is desirable because in some contexts people describe objects as evidence: in a court case, for example, a blood-stained knife might be referred to as a piece of evidence.

Practically speaking, it is advantageous to be able to represent states of partial belief using probability functions whose domains do not necessarily include representions of all possible states of evidence. In many cases it is not possible, let alone useful, to enumerate all possible states of evidence.

On the other hand, the domain-based method is entrenched in scientific working practices and lends itself to the easy explication of certain inductive assumptions. For example, the assumption that evidence of past instances renders future instances more plausible is straightforwardly explicated by the axiom of positive instantial relevance. We shall see below that it is not clear how this assumption should be represented using the calibration-based method.

5.3   Objective Bayesian inductive logic

The first way of relating objective Bayesian epistemology and inductive logic articulates a distinct objective Bayesian form of inductive logic similar to the subjective Bayesian form discussed in chapter 3.

Objective Bayesian inductive logic aims to formalise the inductive assumptions that might be made by agents who are rational according to objective Bayesian epistemology. Just as in the subjective Bayesian case, objective Bayesian inductive logic is distinguished from objective Bayesian epistemology because it has a different goal: it aims to produce useful formal models of inductive reasoning rather than true claims about rational states of partial belief. The general instructions of objective Bayesian inductive logic are as follows:

Objective Bayesian inductive logic
Suppose that $\mathbb{P}$ is the set of probability functions on a certain domain D whose members represent propositions whose plausibility is of interest. In addition, suppose that the relevant evidence in a particular situation determines a set of calibrated probability functions $\mathbb{E}$ and that $\Downarrow \mathbb{E}$ is the set of sufficiently equivocal calibrated probability functions corresponding to $\mathbb{E}$.

Then, for the purposes of formalising the inductive reasoning of a rational agent, use only systems of inductive logic $(D, \mathcal{M})$ satisfying three conditions:

   OBIL1 (Probability): $\mathcal{M} \subseteq \mathbb{P}$

OBIL2 (Calibration): $\mathcal{M} \subseteq \mathbb{E}$

OBIL3 (Equivocation): $\mathcal{M} \subseteq\Downarrow \mathbb{E}$

As in the subjective Bayesian case, objective Bayesian inductive logic is largely determined by its epistemological counterpart. *OBIL*1 excludes sub-probabilistic systems of inductive logic, as befits a form of inductive logic that accepts the probability norm for rational beliefs. *OBIL*2 follows from objective Bayesian epistemology's commitment to the calibration norm, and *OBIL*3 from the equivocation norm. The three stipulations increase in logical strength: *OBIL*3 entails *OBIL*2, which entails *OBIL*1.

This approach seems to match the outline of the objective Bayesian inductive logic presented in (Williamson, 2015, § 5). Williamson construes the norms of objective Bayesian epistemology as providing a semantics for inductive logic: this is essentially the situation if one follows the instructions of objective Bayesian inductive logic. To see why this is the case, consider Williamson's inductive logical schema:

$$\phi_1^{X_1}, \ldots, \phi_k^{X_k} \approx \psi^Y$$

According to Williamson, objective Bayesian inductive logic defines semantic rules for this kind of entailment relationship under a certain interpretation. On this interpretation, $\phi_1, \ldots, \phi_k, \psi$ are members of a set D representing propositions of interest, $X_1, \ldots, X_k$ are sets of probabilities, representing states of evidence about these propositions and Y is a set of probabilities representing possible rational degrees of belief in ψ, given the corresponding evidence. Williamson defines objective Bayesian inductive logic as stipulating that an entailment relationship fitting the above schema obtains if and only if Y contains the value $p_{\Downarrow\mathbb{E}}(\psi)$ of every probability function $p_{\Downarrow\mathbb{E}}$ in the set $\Downarrow \mathbb{E}$, where $\mathbb{E}$ is a set of calibrated probability functions determined by the values in $X_1, \ldots, X_k$. See (Williamson, 2015, § 5.4) for a full discussion of this procedure.

Similarly, the instructions of objective Bayesian inductive logic define semantic rules governing the relationships between sets $\mathbb{E}$ and $\mathcal{M}$, under the interpretation that these sets represent states of evidence and possible rational states of belief. Williamson's formulation differs from mine in two main ways. First, Williamson's presentation specifies how a set $\mathbb{E}$ of calibrated probability functions is to be determined based on a collection $X_1, \ldots, X_k$ of sets of probabilities, whereas my presentation leaves this question open. Second, Williamson's presentation concerns a set Y representing possible rational degrees of belief in a single proposition, whereas my presentation concerns a set $\mathcal{M}$ representing possible rational states of belief in many propositions. Neither of these differences is fundamental.

Objective Bayesian inductive logic, defined according to the above instructions, is distinct from both Carnapian and subjective Bayesian inductive logic. It is

incompatible with Carnapian inductive logic because, like the subjective Bayesian approach, it rejects certain systems of inductive logic—for example, subprobabilistic ones—in all circumstances, for purely philosophical reasons. This amounts to a rejection of Carnap's hard core commitment to tolerating all systems of inductive logic at a philosohical level. In particular, due to its epistemological commitment to the calibration norm, objective Bayesian inductive logic must represent evidence using the calibration-based method rather than the conditional probability method, regardless of the methodological merits of this decision. Carnapian inductive logic may in principle use any method.

Objective Bayesian inductive logic is also incompatible with subjective Bayesian inductive logic, as it uses the calibration-based method to represent evidence, whereas subjective Bayesian inductive logic is committed only to using the conditional probability method.

In the next few sections I compare objective Bayesian inductive logic with Carnapian inductive logic. I argue that objective Bayesian inductive logic has both relative disadvantages and relative advantages. The relative disadvantages are that objective Bayesian inductive logic's commitment to probabilistic necessity is philosophically dubious, and that its commitment to the calibration-based method makes it difficult to represent certain kinds of evidence. On the other hand, objective Bayesian inductive logic leaves less discretion to inductive logicians than Carnapian inductive logic, and can formalise assumptions about physical probabilities that are difficult for Carnapian inductive logic to capture.

5.3.1 Scope for formalising different inductive assumptions

The main way for objective Bayesian inductive logic to represent different inductive assumptions is by varying the set $\mathbb{E}$ of calibrated probability functions so as to represent different states of evidence. Representing different assumptions by adjusting other aspects of objective Bayesian inductive logic is difficult, given its epistemological commitments.

Objective Bayesian epistemology takes probability functions' domains to represent the set of propositions under consideration: this set does not necessarily change depending on the inductive assumption in question. As a result, using different domains to represent different assumptions seems inappropriate, except where different propositions are being represented.

The required level of equivocation also cannot be changed so as to model different inductive assumptions: it is explained at (Landes and Williamson, 2013, p.15) that the main motivation for not formulating the equivocation norm so as to rule out all non-maximally-equivocal probability functions is to deal with unusual constraints. In order to keep in the spirit of objective Bayesian epistemology, the threshold must be set so high that near-maximal equivocality

is required. If this is done, then the set of sufficiently equivocal probability functions will typically be very small.

In contrast, Carnapian inductive logic can represent inductive assumptions by varying all parts of a system of inductive logic: both the set of measure functions and the domain. The set of measure functions can be manipulated by choosing axioms that reflect the assumption that is to be formalised, whether evidential or not. For example, constant exchangeability might be appropriate if, according to the available evidence, a list of observations records successive rolls of a die, in which case the order of observations would impart no information, but inappropriate if the evidence specifies that the list records successive states of the weather, in which case the order of observations would be important. Carnap makes this point at (Carnap, 1971a, p.129). Similarly, different domains can be chosen so as to reflect different kinds of assumption. Carnap discussed this kind of manipulation at (Carnap, 1971a, p.49-52): for some applications, he argued, it is appropriate to formulate evidence in an 'observation' language, whereas for other applications the same evidence can be formulated within a 'theoretical' language.

5.3.2 Evidence about physical probabilities

Objective Bayesian epistemology offers ample instructions as to how the set $\mathbb{E}$ should vary so as to reflect evidence about physical probabilities, or phenomena that behave like them. Objective Bayesian inductive logic is therefore well-equipped to represent inductive assumptions about which degrees of plausibility are appropriate given such evidence. In addition, objective Bayesian inductive logic can easily handle inductive assumptions about degrees of plausibility in light of evidence that fixes the expected values on certain random variables. Both of these forms of evidence directly rule out certain probability functions, and so can naturally be represented as constraints on sets of calibrated probability functions.

This kind of evidence is difficult to depict for inductive logics that rely on the conditional probability method of representing evidence. According to this method, evidence about physical probabilities, or evidence about the expected values of random variables, must be represented as domain-members: this might make the resulting system of inductive logic very unwieldy.

Since it can formalise assumptions about physical probabilities relatively succinctly, objective Bayesian inductive logic therefore has an advantage compared to subjective Bayesian inductive logic, which is committed to using the conditional probability method in all circumstances. It also has an advantage compared to Carnapian inductive logic: although the latter discipline is not committed to using the conditional probability method in all circumstances, it has not yet developed a method that can represent assumptions about physical probability as succinctly as objective Bayesian inductive logic.

However, inductive assumptions that do not have to do with physical probabilities and do not fix the expected values of random variables pose a potential problem for objective Bayesian inductive logic. In particular, it is not immediately obvious how to formalise inductive assumptions of instantial relevance: those according to which instances of a certain kind of observation affect the plausibility of future instances of the same kind. I shall now investigate whether objective Bayesian inductive logic is sufficiently comprehensive to overcome this difficulty.

5.3.3   Evidence about relevance

Consider the case of a nature enthusiast who wants to find out about the habits of the elusive, possibly mythical, Suffolk tiger. This species of tiger is said to roam the countriside in the county of Suffolk in South East England, mostly at night, and always when nobody is looking.

The enthusiast sets up tiger-sensitive cameras in two propitious-seeming locations—one on a hill, and another in a path through a wood—and then returns home and waits for the traps to send images to a computer. Initially, the enthusiast's evidence and background knowledge about the situation are very limited. However, while the enthusiast is at home waiting for images, their friend—a renowned tiger expert—pops in and says that, if the camera in the wood were to spot a tiger, that would make it more plausible that the camera on the hill would spot one as well. In other words, the enthusiast receives evidence according to which one event, namely observing a tiger in the wood, is positively relevant to another, namely observing a tiger on the hill.

Suppose that, in order to benefit from the clarity of formal language, the enthusiast decides to formalise their reasoning in this situation using objective Bayesian inductive logic.

A very simple domain could be used to represent the possibilities, namely the set SL_{tiger} of sentences of a unary predicate language L_{tiger} with one predicate T and two constant symbols a_1 and a_2. Each constant symbol represents a trap-experiment, and the predicate T represents a positive result, i.e. the trap taking a photo of a tiger. L_{tiger} has four state descriptions—$Ta_1 \wedge Ta_2$, $Ta_1 \wedge \neg Ta_2$, $\neg Ta_1 \wedge Ta_2$ and $\neg Ta_1 \wedge \neg Ta_2$—corresponding to the four possible combinations of outcomes of the enthusiast's two experiments.

In addition, a set of measure functions representing possible degrees of plausibility would be required. According to objective Bayesian inductive logic, this should be a subset of the set $\Downarrow \mathbb{E}$ of sufficiently equivocal probability functions that are calibrated with the available evidence. The main problem is to find a set of probability functions $\mathbb{E}_{pre}$ that are calibrated with the enthusiast's evidence before receiving advice from the friendly tiger expert, and another set $\mathbb{E}_{post}$ of functions that are calibrated with the enthusiast's evidence after receiving the new information.
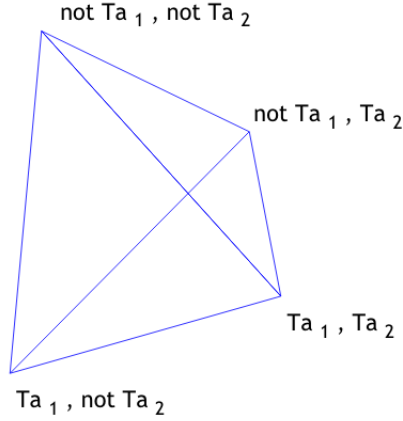
Figure 5.1: The set $\mathbb{P}_4$ of probability functions on SL_{tiger}

Since the enthusiast initially has very limited evidence and background knowledge about the Suffolk tiger, either statistical or qualitative, $\mathbb{E}_{pre}$ should plausibly be the set $\mathbb{P}_4$ of all probability functions whose domains are SL_{tiger}. One might argue that any reasonable person would have good evidence that there are no tigers in Suffolk, given this county's tiger-unfriendly environment, the fact that there are very few tigers of any kind remaining in the world and the absence of any other big cat that is capable of living in a densely populated area without leaving any trace of its existence. I shall assume that the enthusiast wishes to formalise a more open-minded inductive attitude.

This initial set $\mathbb{P}_4$ is represented geometrically in figure 5.1 below.

Each point in the tetrahedron represents a probability function whose domain is SL_{tiger}. Each vertex represents a probability function that gives probability 1 to one state description and probability 0 to all others. Edges represent probability functions that give two state descriptions positive probability, other points on the surface represent functions that give three state descriptions positive probability, and points in the interior allocate positive probability to all state descriptions. The point at the centre of the tetrahedron represents the equivocator function, which gives the same probability $1/4$ to all four state descriptions.

A natural choice for $\mathbb{E}_{post}$ is the set $\mathbf{IR}^+$ of probability functions that satisfy the axiom of positive instantial relevance; that is, those functions p such that $p(Ta_2 \mid Ta_1) > p(Ta_2)$. To see which region of the $\mathbb{P}_4$ tetrahedron $\mathbf{IR}^+$ occupies, consider the independence surface depicted below, which represents the

Figure 5.2: The set of probability functions p such that $p(Ta_2 \mid Ta_1) = p(Ta_2)$

probability functions p such that $p(Ta_2 \mid Ta_1) = p(Ta_2)$. $\mathbf{IR}^+$ occupies the region below this surface.

Unfortunately, objective Bayesian inductive logic cannot straightforwardly choose $\mathbf{IR}^+$ as $\mathbb{E}_{post}$ because $\mathbf{IR}^+$ is non-convex. This can be seen from the diagram: there are line-segments connecting points in $\mathbf{IR}^+$ that pass outside of it.

Objective Bayesian epistemology typically requires sets of calibrated probability functions to be convex because it can be technically difficult to apply the equivocation norm to non-convex sets. Identifying which probability functions within a convex set have the greatest Shannon entropy is, in general, a much simpler problem than in the non-convex case. The former problem is one of convex optimisation whereas the latter is a problem of non-convex optimisation. Boyd and Vandenberghe note at (Boyd and Vandenberghe, 2009, §1.3.1) that "very effective" methods are available for solving problems of convex optimisation, and that " it is reasonable to expect that solving general convex optimization problems will become a technology within a few years", offering no similar assurances for non-convex optimisation problems. Allowing $\mathbb{E}_{post}$ to be non-convex while measuring equivocality using Shannon entropy therefore leads to a methodological problem: finding a set $\Downarrow \mathbb{E}_{post}$ of sufficiently equivocal probability functions becomes more difficult.

Williamson makes a proposal at (Williamson, 2010, p.45)—the convex-hull method— for ensuring that all calibrated sets are convex. According to this method, if one's evidence suggests adopting a non-convex set of probability functions Ev, then one should take as one's set of calibrated probability functions Ev's convex hull $\langle Ev \rangle$. The convex hull of a set of probability functions is the smallest set containing Ev that is closed under convex combinations.

Whatever the merits of the convex hull approach in general, it does not provide

an attractive solution in our case. Following the convex hull approach would suggest taking $\langle \mathbf{IR}^+ \rangle$, the convex hull of $\mathbf{IR}^+$ as the set of calibrated probability functions $\mathbb{E}_{post}$. However, as can be seen from the figure above, this set would contain all of the probability functions in $\mathbb{P}_4$. Following the convex hull approach would therefore make it the case that $\mathbb{E}_{pre} = \mathbb{E}_{post} = \mathbb{P}_4$. This would mean that the evidence provided by the enthusiast's friend had no effect on the enthusiast's epistemic situation, which would be very unintuitive as the friend's evidence seemed informative.

In order to formalise their reasoning using objective Bayesian inductive logic, then, the enthusiast seems to be faced with a choice between two unattractive options: allowing a non-convex set of calibrated probability functions or interpreting the friend's evidence as uninformative.

I shall now consider some possible ways out of this predicament.

Calibrating to a different set

One response to this dilemma is to represent the states of belief that are calibrated with the enthusiast's second evidential state using a set of probability functions other than $\mathbf{IR}^+$. Such an alternative might well be attractive to an objective Bayesian inductive logician in any case, as the appeal of $\mathbf{IR}^+$ depends on taking the conditional probability $p(Ta_2 \mid Ta_1)$ to indicate the degree to which the event represented by Ta_2 is plausible, given evidence according to which the event represented by Ta_1 occurred. In other words, choosing $\mathbf{IR}^+$ takes for granted an instance of the conditional probability approach to representing evidence. Since objective Bayesian epistemology is committed to a method of representing the effect of evidence that is in general inconsistent with the conditional probability approach, an objective Bayesian epistemologist's first instinct might be to lay that this way of representing relevance is the cause of the problem.

There are two obstacles blocking this escape route. First, it is hard to think of an alternative set of probability functions to $\mathbf{IR}^+$ that would represent the enthusiast's second evidential state without appealing to the conditional probability approach. Even if an alternative set were found, this set is not guaranteed to be convex. Second, it is difficult to see why an objective Bayesian should object to taking $p(Ta_2 \mid Ta_1)$ to represent the degree to which a positive result at trap two is plausible, given that trap two took a picture of a tiger. None of the conditions outlined at (Williamson, 2011, p. 73), under which objective Bayesian updating disagrees with the conditional probability approach, seem to apply in this case.

Both of these obstacles would need to be overcome in order for the problem to be resolved by starting with an alternative set of probability functions.

Accept that $\mathbb{E}_{post} = \mathbb{P}_4$

Alternatively, objective Bayesian inductive logic could accept the outcome of the convex-hull method and formalise the enthusiast's second evidential state as the whole of $\mathbb{P}_4$, while interpreting it so as to avoid the counter-intuitive implication that the enthusiast should ignore their friend's evidence.

One way to do this is might be to claim that the friend's evidence has further significance, beyond constraining the choice of $\mathbb{E}_{post}$. Perhaps it constrains the enthusiast's future choice of calibrated probability functions in case they learn the outcome of the first trap experiment before learning about the second. If this were the case then an objective Bayesian could perhaps allow that the enthusiast's second evidential state is represented by the whole of $\mathbb{P}_4$ without having to conclude that the enthusiast has to ignore the new evidence.

This approach faces a general problem. If the friend's evidence affects the enthusiast's reasoning in a way that is not captured in the choice of $\mathbb{E}_{post}$, then objective Bayesian inductive logic does not fully formalise their reasoning.

How significant is this problem?

In sum, while the enthusiast's reasoning might possibly be formalised using objective Bayesian inductive logic, a plausible way to do so has not yet been found.

The problem seems to generalise in at least two ways. First, the particular interpretation, with cameras, tigers and an enthusiast, was inessential: the problem applies to any structurally similar situation, in which evidence must be taken into account according to which one experimental result is probabilistically relevant to another. Secondly, the number of experiments was not important. While the case with two experiments is the only one where the situation can easily be visualised, the set of probability functions satisfying the axiom of positive instantial relevance is also non-convex in cases with greater numbers of experiments.

Objective Bayesian inductive logic therefore faces a challenge: how should it formalise evidence according to which results of a certain kind are relevant to others of the same kind, if not by calibrating to the set of probability functions that satisfy the axiom of positive instantial relevance?

A qualification

Williamson (2013) outlines how objective Bayesian inductive logic can formalise a certain kind of evidence about relevance, namely evidence according to which, if a certain result is frequent in a long sequence of independent and identically distributed observations, then subsequent observations of the same kind should be more plausible.

Williamson's argument can be illustrated using the following example. An agent must decide to what degree they believe the proposition—call it H_{100}—that the result of the hundredth coin toss in a sequence would be heads. The possibilities are heads and tails, so the agent must choose a function within the set $\mathbb{P}_2$ of probability functions with a two-member domain. In the absence of any evidence one way or the other the calibrated set is the whole of $\mathbb{P}_2$, and the most equivocal choice assigns probability $1/2$ to H_{100}.

However, suppose the evidence includes the results of the previous 99 tosses—suppose they were all heads—together with the information that each toss should be regarded as probabilistically independent and as assigning the same probability to heads as the probability that should be assigned to H_{100}. Classical statistical methods teach us that, in these circumstances, it is very plausible that H_{100} should be assigned a probability greater than $6/10$. If the 'ruling out' proposition that H_{100} should be assigned a probability greater than $6/10$ is plausible enough in light of the results, then a rational agent should arguably incorporate this proposition into their evidence. The set of calibrated probability functions is now not the whole of $\mathbb{P}_2$, but rather the subset of probability functions that assign H_{100} probability greater that $6/10$. The most equivocal choice now assigns probability $6/10$ to H_{100}: the combined result of the instances, together with the assumption of independent and identical objective probabilities, was to increase the probability of H_{100}.

This example seems to show that, given the right background conditions, objective Bayesian inductive logic can accommodate evidence about positive relevance. However, for several reasons, only a limited array of evidence about relevance can be represented in this way. First, the background conditions need to obtain - if one's evidence suggests that there should be positive relevance in a case that does not feature a long sequence of independent and identically distributed events, then this evidence cannot be formalised using Williamson's method. Second, the degree of positive relevance depends on the exact threshold at which the ruling out proposition becomes plausible enough to be incorporated as evidence. Whether all rational agents have 'Lockean thresholds' is philosophically controversial—see, for example Hawthorne (2009)—and identifying a suitable threshold for a particular application would surely be difficult. Even if a threshold is found in a particular case, the degree of positive relevance that it and the available statistical evidence imply may not match the degree of positive relevance suggested by non-statistical evidence about relevance. Finally, it is not clear that this method could be extended for the case of evidence about negative relevance.

5.3.4   Conclusion

To sum up the preceding discussion, it is possible to construct a distinct objective Bayesian form of inductive logic that draws on objective Bayesian epistemology. This objective Bayesian inductive logic offers a good framework for

formalising inductive reasoning that incorporates evidence about physical probabilities and evidence that fixes the expected values of random variables. On the other hand, objective Bayesian inductive logic faces a challenge as to how it should formalise evidence about relevance. So far, this challenge has not been met in full.

I shall now briefly compare objective Bayesian inductive logic with Carnapian inductive logic.

At a philosophical level, Carnapian inductive logic seems to be preferable to objective Bayesian inductive logic. Whereas Carnapian inductive logic makes only comparatively uncontroversial philosophical commitments, objective Bayesian inductive logic is committed to probabilistic necessity, which is epistemologically dubious.

From a methodological point of view, objective Bayesian epistemology is far more specific than Carnapian inductive logic. Whereas Carnapian inductive logic affords much discretion to the practising inductive logician—for example, which axioms to impose, which method to represent evidence with—objective Bayesian inductive logic makes these decisions automatically. As a result, objective Bayesian inductive logic seems to be preferable to Carnapian inductive logic for applications that involve the kind of evidence that it specialises in. On the other hand, Carnapian inductive logic can represent inductive assumptions that are difficult for objective Bayesian inductive logic to formalise, such as assumptions involving evidence about instantial relevance.

In order to determine whether Carnapian inductive logic or objective Bayesian inductive logic should be used for a given application, it must be determined what kind of evidence is to be formalised. If the application involves only formalising evidence about physical probabilities, then objective Bayesian inductive logic may well be preferable to Carnapian inductive logic due to its specificity. On the other hand, if the application involves evidence about instantial relevance, then Carnapian inductive logic is likely to be preferable.

Whether Carnapian and objective Bayesian inductive logic turn out to be viable will depend on whether there are sufficiently many, sufficiently important, applications of the kind that it is most suited for. This is difficult to predict in advance; nonetheless, in chapter 6 I will attempt to outline an area of application for Carnapian inductive logic—to formalising assumptions that arise during statistical investigations—which is both important and to which other forms of inductive logic do not seem apt. Whether there are niches in the scientific landscape that are similarly suitable for objective Bayesian inductive logic remains to be seen.

5.4 The assimilative approach

Rather than attempting to find a distinctive objective Bayesian inductive logic, it is also possible to treat the objective Bayesian norms as inductive assumptions and attempts to formalise them within Carnapian inductive logic. In this section I examine this 'assimilative' approach to the problem of relating objective Bayesian epistemology and inductive logic.

I attempt to describe, in general terms, what a system of inductive logic that formalises the objective Bayesian norms would look like and to judge the viability of the assimilative approach on that basis.

Before doing so, it is necessary to address an objection to the very idea of formalising the objective Bayesian norms within Carnapian inductive logic. A proponent of objective Bayesian epistemology might, at first glance, consider the assimilative approach to be fundamentally at odds with the spirit of objective Bayesian epistemology. Epistemological principles, they might contend, are quite different from inductive assumptions. An inductive assumption is speculative—depending on the facts, it might be wise or unwise—whereas an epistemological principle is not: acting in accordance with a genuine principle of rationality is wise regardless of the facts. By construing the norms of objective Bayesian epistemology as inductive assumptions, the proponent might continue, the assimilative approach risks glossing over any debate about whether the norms are genuine rationality requirements. Even worse, the assimilative approach might even seem to pre-judge this debate, presupposing that the objective Bayesian norms are not rationality requirements but only mere assumptions.

I can think of two responses to this objection. First, it is not clear that objective Bayesian epistemologists typically conceive of their norms as absolutely exceptionless principles. For example, Williamson acknowledges in his discussion of the probability norm at (Williamson, 2010, p. 38) that that the probability norm need not constrain the partial beliefs of agents who find themselves in unusual betting scenarios. Similarly, objective Bayesian epistemologists might allow that, in certain unusual circumstances, agents might rationally contravene the other norms without threatening those norms' statuses.

Thus it is not clear that practising epistemologists hold to the kind of sharp distinction between exceptionless principles and fact-dependent assumptions that the objection envisages. As a result construing its norms as assumptions perhaps does less of an injustice to the objective Bayesian epistemology than might be thought at first glance.

Secondly, and more importantly, adopting the assimilative approach need not involve the epistemological presuppositions that the objection implies. Carnapian inductive logic does not aim to settle epistemological questions, but rather to avoid them by focusing on the practical, rather than epistemological, merits of formalising different views. Formalising the objective Bayesian norms within

Carnapian inductive logic does not necessarily commit one to the view that they are not genuine rationality principle. One need only decide that formalising the norms is useful: it is possible to maintain this position at the same time as thinking that the norms are genuine epistemological principles.

To take a similar case, both someone who thinks that the future tends to resemble the past and an inductive sceptic who doubts that this is the case might have good pragmatic reasons to formalise the assumption of positive instantial relevance within Carnapian inductive logic. Analogously, both people who think that the objective Bayesian norms are universal principles and those who doubt this may reasonably adopt the assimilative approach.

5.4.1   Objective Bayesian epistemology as a system of inductive logic

In order to carry out the assimilative approach, a system of inductive logic—that is, a domain and a set of permitted measure functions—must be found that formalises the norms of objective Bayesian epistemology.

The measure functions that may represent rational states of belief, according to objective Bayesian epistemology, will be characterised by axioms corresponding to the norms set out in section 5.2.1 above. Following the normal approach of Carnapian inductive logic, the effect of taking evidence into account will be represented using the conditional probability method.

The main challenge will be to find a way, consistent with the conditional probability approach, to formulate the calibration norm. This norm conceives of a body of evidence as a collection of constraints as to which measure functions can be adopted, rather than as a proposition that can be conditionalised upon.

Fortunately, by developing an argument from Skyrms (1985) it is possible to find a method for formalising the calibration norm. Skyrms argues that objective Bayesian updating—the process of repeatedly applying the objective Bayesian norms under changes of evidence—can be analysed as a special case of the application of the rule of conditionalisation. Skyrms's key idea is to expand the domains of the conditionalising probability functions so that they include members that represent what he calls "propositions about" physical probabilities.

The same idea can work for our purposes: in order to describe the calibration norm as an axiom, we must consider measure functions with rich domains whose members can represent not only propositions about possible states of affairs, but also propositions about possible evidential states. Call such a domain D_+. I shall now describe how such a domain can be constructed.

In order to represent all relevant propositions about states of the world, the new domain D_+ must contain a sub-domain whose members represent such states. Within objective Bayesian epistemology the most basic approach, followed in Williamson (2010) and Landes and Williamson (2013), is to presume that the

objects of rational agents' beliefs can be represented by members of a finite propositional domain D_{prop}. I will follow this approach in what follows: the richer domain D_+ will be taken to contain a propositional domain D_{prop} whose members will be taken to represent propositions about possible states of affairs.

As a result of starting in this way with D_{prop}, rather than a richer domain, the approach sketched here will not formalise the more sophisticated form of objective Bayesian epistemology explored in Landes and Williamson (2015) where predicate domains are considered. Nonetheless, a similar analysis ought to apply to such cases with appropriate modifications.

In addition to containing D_{prop}, D_+ must contain another sub-domain D_{Ev} whose members represent possible states of evidence, construed as sets of calibrated probability functions over D_{prop}. Since any set of probability functions on D_{prop} can in principle be a calibrated set, it seems like D_{Ev} must be $\mathcal{P}(\mathbb{P}_{D_{prop}})$, the powerset of the set of all probability functions over D_{prop}.

Beyond the fact that it must contain both both D_{prop} and D_{Ev}, I will not explore exactly how D_+ ought to be constructed. One way to construct D_+ so as to contain both D_{prop} and D_{Ev} would be to specify that $D_+ = D_{prop} \cup D_{Ev}$, setting D_+ as the union of these two sets. Many issues need to be addressed in order to determine whether this is the best choice. For example, would it be preferable for D_+, unlike $D_+ = D_{prop} \cup D_{Ev}$, be a boolean algebra, and therefore closed under union and intersection? If so, how should we interpret members like $X \cup \mathbb{E}$, where $X \in D_{prop}$ and $\mathbb{E} \in D_{Ev} = \mathcal{P}(\mathbb{P}_{D_{prop}})$? While such issues would need to be addressed in order to carry out the assimilative approach fully, they do not bear directly on this programmatic discussion, for which purposes we may assume that a suitable domain D_+ can be found.

Suppose, therefore, that a suitable domain D_+ is found that contains both D_{prop} and D_{Ev}. We can now think about how to reconstruct the norms of objective Bayesian epistemology as axioms for measure functions whose domains are D_+.

The probability norm can be reproduced as an axiom requiring measure functions on D_+ to have restrictions to D_{prop} that are probability functions.

In order to reconstruct the calibration norm, some way must be found to distinguish the set of probability functions that represents the state of evidence associated with a given state of belief. One way to do this would be to stipulate as a preliminary axiom that each measure function over D_+ should assign the value 1 to exactly one member of D_{Ev}—taken to represent its corresponding state of evidence—while assigning the value 0 to all the other sets of probability functions in D_{Ev}. The calibration norm may then be reconstructed as the requirement that, if a measure function m assigns the value 1 to a set of probability functions $\mathbb{E}$, then $m_{\restriction D_{prop}}$, the restriction of m to D_{prop}, must be a member of $\mathbb{E}$.

Finally, the equivocation norm can be reconstructed along similar lines to the calibration norm. It must stipulate that, if m assigns the value 1 to $\mathbb{E}$, then

$m_{\restriction D_{prop}}$ must be a member of the set $\Downarrow \mathbb{E}$.

To sum up, given a domain D_+ containing a set of propositions D_{prop} and a set D_{Ev} of sets of probability functions over that domain, the norms of objective Bayesian epistemology are plausibly captured by the following axioms:

> **Objective Bayesian axioms**
> Allow only measure functions $m : D_+ \to \mathbb{R}$ such that, for all $\mathbb{E} \in D_{Ev}$:
>
>> **Probabilism:** $m_{\restriction D_{prop}} \in \mathbb{P}_{D_+}$.
>>
>> **Pre-calibration:** m assigns the value 1 to exactly one member of D_{Ev}, and assigns the value 0 to all other members of D_{Ev}.
>>
>> **Calibration:** If $m(\mathbb{E}) = 1$, then $m_{\restriction D_{prop}} \in \mathbb{E}$.
>>
>> **Equivocation:** If $m(\mathbb{E}) = 1$, then $m_{\restriction D_{prop}} \in \Downarrow \mathbb{E}$.

Suppose that these axioms are satisfied by a set of measure functions $\mathcal{M}_{OBE}$. The system of inductive logic that explicates the norms of objective Bayesian epistemology within Carnapian inductive logic is then $(D_+, \mathcal{M}_{OBE})$.

Since it seems possible in principle to construct systems like $(D_+, \mathcal{M}_{OBE})$, it therefore appears possible to formulate the norms of objective Bayesian epistemology within Carnapian inductive logic.

5.4.2   Advantages of the assimilative approach

What could be the benefit of formalising the norms of objective Bayesian epistemology within Carnapian inductive logic? I can think of two. First, formalising the norms makes it possible to compare objective Bayesian systems of inductive logic like $(D_+, \mathcal{M}_{OBE})$ with other systems of inductive logic. This may help to inform philosophical discussions about the relationship between objective Bayesian epistemology and alternative epistemological positions. Secondly, by spelling out the main claims of objective Bayesian epistemology fully formally, the assimilative approach can clarify its substantive claims about rationality, making it easier to assess their plausibility.

Comparison

If the assimilative approach is carried out, it will be possible to compare objective Bayesian systems of inductive logic like $(D_+, \mathcal{M}_{OBE})$ with other systems of inductive logic. In particular, it will be possible to compare objective Bayesian systems with systems representing alternative epistemological positions. This may help to clarify the differences between these positions more than would be possible without such a formalisation, and also allows pragmatic argumentation

to be introduced into debates about objective Bayesian epistemology in a way that would not be available otherwise.

Specifically, thanks to the assimilative approach, objective Bayesian epistemology can be compared with alternative positions based on the pragmatic merits of the systems of inductive logic that formalise them. Systems of inductive logic are easier to compare on a pragmatic basis than informal epistemological claims, firstly because they are fully explicit and secondly because many areas of science are already in the business of comparing similar formal structures, such as statistical models. As we shall see in the next chapter, there are strong parallels between statistical models and systems of inductive logic. As a result, the same, or similar, methods that statisticians use to compare the usefulness of statistical models can potentially be used to compare objective Bayesian systems of inductive logic with alternatives that formalise rival positions.

Comparing the pragmatic merits of different systems of inductive logic will not definitively settle debates about the relative merits of their corresponding epistemological doctrines. At best it will show that one doctrine is more germane to useful formalisation than the other. However, such pragmatic considerations can still be relevant to epistemological debates. Often such debates can only be resolved by choosing between incompatible yet equally empirically adequate intuitions; such decisions can be very difficult to justify one way or the other. In these circumstances, any difference between two epistemological positions, even if it is 'merely' pragmatic, can be useful as a way of breaking the deadlock.

As an example of the potential benefit of connecting debates about objective Bayesian epistemology with pragmatic discussions using the assimilative approach, consider the objective Bayesian conception of evidence. According to this conception, states of evidence should be thought about as determining sets of calibrated probability functions, to be taken into account using the calibration method. Alternatively, according to the conditional conception, states of evidence should not be thought of as sets of probability functions, but rather as propositions to be taken into account using the conditional probability method.

There has been much philosophical debate about how the objective Bayesian and conditional conceptions of evidence compare philosophically. For example, Skyrms (1987) argues that the objective Bayesian conception captures taking evidence into account in a 'supposing' kind of way, whereas the conditional conception describes taking evidence into account in an 'updating' kind of way. Shimony (1985) argues that the conditional conception is generally preferable on philosophical grounds. Williamson (2011) gives a philosophical defence of the objective Bayesian conception.

The assimilative approach can introduce a new way to resolve this debate by comparison of the practical merits of each conception's inductive logical counterpart.

For a given application of inductive logic, the objective Bayesian conception of evidence can be represented by the specification of a domain with a structure

like that of D_+, together with the 'pre-calibration' and calibration axioms. Alternatively, the conditional conception could be represented by the specification of a domain representing possible evidence propositions, together with axioms imposing conditions on conditional measure functions, such as, for example, axioms of instantial relevance or Johnson's sufficientness postulate.

If systems of inductive logic in which the objective Bayesian framing assumption is represented, like $(D_+, \mathcal{M}_{OBE})$, turn out to be more useful in practical applications than systems of inductive logic that represent the conditional conception, then this would constitute an argument in favour of the objective Bayesian conception of evidence, and vice versa.

For applications where the available evidence merely reveals the truth of certain propositions, objective Bayesian systems of inductive logic like $(D_+, \mathcal{M}_{OBE})$ are likely to be unnecessarily complicated. Inductive assumptions about reasoning based on this kind of evidence could be formalised more simply using a conditional-probability-based system of inductive logic with a comparatively simple propositional domain. On the other hand, objective Bayesian systems might be more appropriate for formalising assumptions about how to respond to evidence that reveals facts about physical probabilities. For such purposes, systems like $(D_+, \mathcal{M}_{OBE})$, although complicated, might be more useful than any purely conditional-probability-based system.

Claims about rationality

A second way in which the assimilative approach can be helpful is to clarify exactly what objective Bayesian epistemology claims about rational states of belief.

For example, consider the calibration norm. Stated informally, the calibration norm might seem like a very strong constraint on rational belief, stipulating that rational agents treat their evidence in a certain way. However, when the objective norms are formalised as above, a different picture emerges.

Rather than the calibration axiom itself, it is the choice of the domain D_+, together with the pre-calibration axiom, which do the most work towards representing how rational agents conceive of their evidence. As we saw in the previous subsection on comparison, there are alternatives to the package of D_+ and the pre-calibration axiom, such as the use of a standard propositional domain whose members represent possible of states of evidence, as suggested by the conditional-probability conception.

However, after the choice to represent evidence using D_+ and the pre-calibration axiom have been made, it is very difficult to think of plausible alternatives to the calibration axiom. According to the interpretation suggested by the choice of D_+ and pre-calibration, measure functions that the calibration axiom excludes represent states of belief that are ruled out by the available evidence. It would

be very difficult to argue that a rational agent should respond to their evidence by contradicting it in this way.

Formalising the objective Bayesian norms therefore seems to make clear an important point: it is not the calibration norm itself that most strongly expresses objective Bayesian epistomogy's claims about how rational agents respond to evidence. Rather, it is the collection of background assumptions according to which rational agents' evidence determines a set of calibrated probability functions that does most of the work.

This observation may help to improve the debate on the objective Bayesian conception of evidence.

5.4.3 Problem: unconnectedness

One problem with the assimilative approach is that systems of inductive logic like $(D_+, \mathcal{M}_{OBE})$, and in particular domains like D_+, are very different from other systems of inductive logic and domains that feature in Carnapian inductive logic. At a formal level, we have seen that D_+ will not necessarily be a boolean algebra. If D_+ is not a boolean algebra, then it will be difficult to apply many familiar axioms which presuppose this condition. For example, the notion of a probability function presupposes a domain that is a boolean algebra. Another formal difference from other Carnapian domains is that, since D_+ has to contain an uncountably infinite set of probability functions, it too must be uncountably infinite.

In addition, even if they have the correct formal structure, domains like D_+ must be interpreted differently from the domains that feature in other systems. Whereas the orthodox approach is to interpret the domains of measure functions as representing possible states of affairs, some members of domains like D_+ must be interpreted as representing states of affairs, while others represent states of evidence. Due to this different interpretation, axioms that explicate interesting inductive assumptions under the assumption that domain-members represent only states of affairs may not do so in the case of domains like D_+.

These differences seem to show that the systems of inductive logic produced by the assimilative approach lack the quality that Carnap called 'exactness'. An explicatum is 'exact', Carnap explains at (Carnap, 1950b, p.7), if it is part of "a well connected system of scientific concepts". Although the term that Carnap used to describe this quality is perhaps somewhat misleading—I will instead use 'well-connectedness'—he was surely right to point out that it is desirable for an explicatum to connect neatly with other related concepts.

Unfortunately, thanks to its unusual domain, $(D_+, \mathcal{M}_{OBE})$ is not very well-connected in comparison to more familiar systems of inductive logic.

For several reasons, I think that the comparative unconnectedness of the systems of inductive logic that it produces is not a decisive objection to the assimilative

approach.

For one thing, any lack of well-connectedness must be weighed against the other virtues of the explicata that the assimilative approach produces, in order to determine whether it is worthwhile, all things considered. Systems of inductive logic like $(D_+, \mathcal{M}_{OBE})$ may turn out to be useful despite being poorly connected.

Secondly, while $(D_+, \mathcal{M}_{OBE})$ may be poorly connected in comparison with other systems of inductive logic, it is arguably still fairly well-connected in an absolute sense. D_+ and $\mathcal{M}_{OBE}$ are unremarkable sets whose formal properties can be mathematically investigated. Moreover, the axioms that define $\mathcal{M}_{OBE}$ are well-connected with the norms that they formalise: these have been comprehensively studied at both a philosophical, scientific and formal level.

Finally, the assimilative approach may improve its connectedness with time. Relatively poor connectedness is to be expected when a research programme begins to produce a new kind of explicatum. Improving the connectedness of the assimilative approach's explicata is a natural next agenda item.

5.4.4   Assessment of the assimilative approach

To conclude, the norms of objective Bayesian epistemology, as presented in Williamson (2010), can in principle be formalised by a system of inductive logic along the lines of $(D_+, \mathcal{M}_{OBE})$. This formalisation is useful as it provides a way to improve debates about objective Bayesian epistemology, by allowing the introduction of otherwise unavailable pragmatic considerations and by clarifying the nature of its claims about rational belief. While the systems of inductive logic that formalise the objective Bayesian norms are somewhat poorly-connected in comparison with those that feature in the rest of Carnapian inductive logic, this is not a decisive objection, all things considered.

It is interesting to note how well the assimilative approach fits in with Carnap's overall approach to philosophical issues.

By attempting to produce a framework within which objective Bayesian epistemology can be compared with alternative epistemological positions, the assimilative approach has a similar goal to Carnap's first major work, the 'Logical Structure of the World', Carnap (1928). According to recent scholarship, as summarised in Leitgeb (2011), Carnap's aim in this work was to produce a framework within which different meta-ontological positions, such as realism and idealism, could be formalised and compared. Carnap's sought to replace a metaphysical debate over the true nature of the universe, which he saw as intractable, with a practical debate over whether to use a 'realistic language' or an 'idealistic language'. As I argue above, the assimilative approach may be used to transform epistemological debates in a similar way: instead of debating whether objective Bayesian norms, or others, genuinely describe rationality, the

assmilative approach allows a practical discussion to proceed as to whether to represent rationality using an objective Bayesian formal language or another.

In addition, the assimilative approach has strong parallels with the strategy that Carnap pursues in Carnap (1967b). In this late essay, Carnap addresses what he calls the problem of uncertain evidence: that is, the problem of formalising the inductive assumptions of rational agents who are not perfectly certain of the outcomes of certain experiments. In order to solve this problem Carnap appeals to systems of inductive logic with domains whose basic elements included both ordinary atomic sentences with the form $P_1 a_i \wedge \ldots \wedge P_k a_i$ and also tuples of numbers with the form $E'_i = (u_{i1} \ldots u_{ik})$, which he called 'evidential quasi-propositions'. Evidential quasi-propositions were intended to represent profiles of credences in each of the k possible outcome of experiment i. The inductive assumptions of rational agents who are unsure of the outcomes of experiments, Carnap proposed, can usefully be represented by axioms for conditional measure functions with this kind of domain.

The assimilative approach can be thought of as an attempt to formalise the assumptions of agents who receive another kind of unconventional evidence: evidence that changes the space of possible states of belief rather than evidence that changes their belief in various experimental outcomes. Just as, in order to describe the latter kind of evidence, Carnap considered measure functions whose domains include evidential quasi-propositions, so the assimilative approach considers measure functions whose domains include sets of probability functions to describe the former.

It seems, then, that at a broad conceptual level and a narrow technical level the assimilative approach is in tune with both Carnap's early and late philosophy.

5.5 Conclusion

I have argued that there are at least two viable ways of relating objective Bayesian epistemology and Carnapian inductive logic. One constructs a distinct objective Bayesian form of inductive logic, which is useful for representing inductive assumptions arising from evidence about physical probabilities, but has a problem representing evidence about relevance. The other formalises the norms of objective Bayesian epistemology within Carnapian inductive logic. This makes it possible to apply pragmatic considerations to debates about objective Bayesian epistemology and, potentially, to clarify the claims of objective Bayesian epistemology.

The two approaches are very different. Whereas the assimilative approach is consonant with both Carnapian inductive logic and Carnap's broader approach to philosophy, objective Bayesian inductive logic departs fundamentally from both. It disagrees with Carnapian inductive logic's commitment to tolerance, arguing for the use of particular systems of inductive logic on philosophical,

rather than practical, grounds. Objective Bayesian inductive logic also takes a different overall philosophical approach: rather than aiming to produce a general framework for formalising all kinds of inductive assumption, it focuses only on formalising inductive assumptions that might be made by rational agents.

Despite these differences, I think that the two approaches should not be seen as opposed to each other. They may even be complementary: the assimilative approach can aid objective Bayesian inductive logic by clarifying the relationship of its underlying norms to similar alternatives. In the other direction, objective Bayesian inductive logic can assist the assimilative approach by providing philosophically and methodologically well-motivated principles for it to formalise.

The preceding discussion has hopefully shown that Carnapian inductive logic has the potential to coexist within the same conceptual habitat as forms of inductive logic based on objective Bayesian epistemology, and even to develop symbiosis with them. Thus the case for substantively reappraising Carnapian inductive logic is surely strengthened by considering of this 'rival' research programme.

Chapter 6

An application

6.1   Introduction

In this chapter I present an application of Carnapian inductive logic to a contemporary philosophical debate. The central question is this: how should the process of choosing between statistical models be conceived philosophically? In Gelman and Shalizi (2012), two practising statisticians compare some philosophical accounts of statistical model-choice: the 'received view of Bayesian inference' and a falsificationist account. They argue that the latter is preferable, as the former is incompatible with certain indispensable features of real-world statistical methodology. I largely agree with these arguments, but think that the falsificationist account is also problematic. A third philosophical account of statistical model choice, informed by Carnapian inductive logic, does a much better job. This application shows that Carnapian inductive logic is relevant to modern philosophical debates, adding further weight to my previous arguments that it merits substantive reappraisal.

The structure of this chapter is as follows. First, I introduce the problem of describing statistical model-choice and discuss its importance. Next I summarise Gelman and Shalizi's arguments against the 'received view' and in favour of the falsificationist account. I then introduce the Carnapian account of statistical model choice, according to which statisticians should treat statistical models analogously to how Carnapian inductive logicians treat systems of inductive logic. I explain this analogy in detail, showing that some statistical models that feature prominently in Gelman and Shalizi (2012) are formally equivalent to systems of inductive logic, and that good reasons for choosing between systems of inductive logic can also be good reasons for choosing between statistical models. I argue that the Carnapian account should be preferred to the falsificationist account for two reasons. First, it is more technically fruitful, suggesting interesting lines of research. Secondly, Carnap's overall philosophical outlook is in much

better harmony with Gelman and Shalizi's account of statistical practice than the philosophy of Karl Popper, which underpins the falsificationist account. I conclude that the Carnapian account is preferable to the falsificationist account and the received view, and finish by discussing the significance of this argument.

6.2   Statistical model-choice

The central concept of this chapter is the philosophical account of statistical model-choice. Gelman and Shalizi (2012) criticises one such account, which it calls 'the conventional philosophy of Bayesian statistics' and proposes a different account inspired by falsificationist philosophy of science. Below I introduce another philosophical account inspired by Carnapian inductive logic and argue that it is better than the falsificationist one.

Philosophical accounts of statistical model-choice are succinct stipulations indicating how statisticians ideally ought to conduct investigations involving statistical models.

Philosophical accounts of scientific methodology are important because they can influence scientific research. As Gelman and Shalizi put it,

> . . . even those [scientists] who believe themselves quite exempt from any philosophical influences are usually the slaves of some defunct methodologist."
> (Gelman and Shalizi, 2012, p.31)

6.2.1   Statistical models

Citing 'currently accepted theories', Mccullagh defines statistical models in general as follows:

> . . . a statistical model is a set of probability distributions on the sample space $\mathcal{S}$.
> (McCullagh, 2002, p. 1225)

A sample space is a set containing a mathematical representation of every possible outcome of a sample of observations. Typically possible sample results are represented by profiles of values of random variables defined on an underlying set Ω of 'states' or 'possible worlds'.

6.2.2   Gelman and Shalizi's arguments

Criticism of the 'conventional philosophy of Bayesian statistics'   Gelman and Shalizi aim to counteract the influence of what they call 'the received view of Bayesian inference', a philosophical account of statistical model-choice

that, they claim, has had a negative effect on statistical research. According to the received view, Gelman and Shalizi write,

> Anything not contained in the posterior distribution $p(\theta \mid y)$ is simply irrelevant...
> (Gelman and Shalizi, 2012, p.9)

I therefore consider the following, slightly more general, stipulation to be a key tenet of the received view:

RV All desiderata that are relevant in a statistical investigation should be represented formally in a statistical model. Other factors should be disregarded.

This stipulation does not amount to a fully-fledged philosophy of statistics and therefore should not be seeen as encapsulating the received view, which must include other stipulations: perhaps that models should be chosen so as to fit given data. Nonetheless, since it is where Gelman and Shalizi focus their criticism, this is the only aspect of the received view that we need to consider.

Gelman and Shalizi argue that **RV** is incompatible with certain facts about statistical research as it goes on in the real-world, as there are important uses in statistical investigations for knowledge that is not represented formally, let alone within a statistical model.

According to Gelman and Shalizi, **RV** is incompatible with certain facts about statistical research as it goes on in the real-world. In practice, they claim at (Gelman and Shalizi, 2012, p.18), properties of statistical models that are unrelated to their performance at representing knowledge are often important, such as mathematical and conceptual tractability.

Similarly, Gelman and Shalizi claim that statistical investigations should sometimes make use of knowledge that is not represented formally. In (Gelman and Shalizi, 2012, §3) they argue that it is practically impossible to represent all the assumptions that might be entertained during the course of an investigation in the form required for a Bayesian statistical analysis. In (Gelman and Shalizi, 2012, §4), Gelman and Shalizi claim that the way statisticians test statistical models against empirical data typically incorporates knowledge that is not represented formally. They cite as an example the knowledge that enables statisticans to recognise whether discrepancies between real data and data simulated using a fitted model are systematic and important or patternless and safe to ignore.

The falsificationist account   Gelman and Shalizi propose an alternative philosophical account of statisticians' model-selection choices that is inspired by falsificationism, a prominent approach to the philosophy of science. Given Gelman and Shalizi's use of the term, I call this philosophical account 'falsificationist'.

F A candidate statistical model should first be chosen for consideration using the statistician's judgement. The model should then be confronted with data and either rejected or cautiously accepted depending on how well it is found to resemble the data source.

Gelman and Shalizi think that the way in which resemblance between models and data sources should be measured, as well as the level of non-resemblance required for rejection, should depend on the particular circumstances of the relevant investigation. They write:

> . . . the hypothesis linking mathematical models to empirical data is not that the data-generating process is exactly isomorphic to the model, but that the data source resembles the model closely enough, in the respects which matter to us, that reasoning based on the model will be reliable.
> Gelman and Shalizi (2012, p.20)

Gelman and Shalizi's principled abstention from specifying exactly how 're-semblance' between statistical models and data sources should be measured distinguishes their account from the received view. According to **RV** only re-semblance that is represented formally in a statistical model can be relevant in a statistical investigation, whereas the falsificationist account allows tests of resemblance that do not have this property.

The key feature of resemblance between models and data, according to Gelman and Shalizi, seems to be that, if there is resemblance, then the assumption that "reasoning based on the model will be reliable" is justified. This desideratum encompasses standard tests of model-fit to the extent that reasoning based on poorly-fitting models is unreliable.

Below I take Gelman and Shalizi's claims about statistical research in practice for granted. Specifically, I assume that working statisticians often have reasons other than knowledge-representation for choosing between statistical models, as well as uses for knowledge that is not represented in the form of a statistical model. In addition I assume that statisticians do and should choose models in roughly the way suggested by Gelman and Shalizi's account **F**, cautiously adopting models after testing them for reasoning-justifying resemblance to data in the situation- and priority-sensitive way that Gelman and Shalizi outline. These assumptions allow me to focus on other aspects of Gelman and Shalizi's account that I find controversial.

Despite the fact that I do not dispute its empirical adequacy, I argue below that Gelman and Shalizi's falsificationist philosophical account should be rejected in favour of an alternative Carnapian one.

In the next section I outline this account before claiming that it is more fruitful than the falsificationist one and also in better harmony with Gelman and Shalizi's arguments.

6.3 The Carnapian account

Many statistical models are formally equivalent to systems of inductive logic, and, I argue in this section, can plausibly be evaluated according to the same criteria. Zabell (2011) makes essentially this point, arguing that statistical research can be improved by characterising families of statistical models according to qualitative conditions discussed within Carnapian inductive logic, such as constant exchangeability and Johnson's sufficientness postulate. The 'Carnapian' account of statistical model-choice that this view suggests can be summarised as follows:

> **C** Statisticians ought to choose statistical models that correspond to systems of inductive logic that are good explications of whichever inductive assumptions they wish to entertain.

I shall now make this connection explicit by showing exactly how statistical models correspond to systems of inductive logic.

Statistical models as systems of inductive logic

As McCullagh points out in the quotation above, statistical models generally consist of a sample space and a set of probability functions. Since sample spaces can be thought of as inductive logical domains, and sets of probability functions are sets of measure functions, there is a trivial sense in which statistical models are systems of inductive logic. However, I shall now argue that there are deeper parallels: in fact, many domains of Carnapian systems of inductive logic can naturally be interpreted as sample spaces, and many sample spaces that are important in statistics can be interpreted in terms of Carnapian inductive logic.

First, consider a unary predicate language $L^{1,n}$ with one predicate R, the usual connectives, no quantifiers and n constants $a_1, a_2, \ldots, a_n$. $L^{1,n}$ has two atoms, $R(x)$ and $\neg R(x)$. The set of equivalence classes of $L^{1,n}$'s sentences can naturally be associated with the sample space generated by the two-valed random variables $X_1, X_2, \ldots, X_n$ defined on the underlying state space $\Omega = \{\omega_1, \omega_2, \ldots, \omega_{2^n}\}$. More generally, if $L^{1,n}$ has q unary predicates, and therefore 2^q atoms, its non-equivalent sentences can be associated with a state space generated by 2^q-valued random variables.

Working in the other direction, state spaces generated by finite-valued random variables can be associated with a set of sentences as follows.

Suppose that a sample space $\mathcal{S}$ is generated by random variables $X_1, \ldots, X_n$, $Y_1, \ldots, Y_n$ and that the random variables from the X category each have v_X possible values, while those in the Y category have v_Y. The required set of sentences will be generated by a unary predicate language $L^{q,n}$ with n constants and $q = v_X + v_Y$ predicates $P_1^X, \ldots, P_{v_X}^X, P_1^Y, \ldots, P_{v_Y}^Y$. This language has one

predicate corresponding to each possible value of each category of random variable. The set $SL^{q,n}$ of sentences generated by $L^{q,n}$ has a member corresponding to every possible sample in $\mathcal{S}$. For example, the sentence $P_{50}^X(a_1) \wedge P_{10}^Y(a_1)$ corresponds to the possibility $(X_1 = 50, Y_1 = 10)$. However, $SL^{q,n}$ also contains some sentences that do not represent possible samples, and so need to be excluded from consideration. For example, $P_{50}^X(a_1) \wedge P_{49}^X(a_1)$ must be excluded as X_1 cannot have the value 49 and 50 at the same time. Fortunately, Carnap demonstrated how to exclude these sentences at (Carnap, 1971a, p. 81). Predicates must be segregated into 'families' corresponding to the categories of random variables that they represent, and it must be stipulated by means of a basic assumption that every constant instantiates exactly one predicate from each family. In our example there will be two families, one containing $P_1^X, \ldots P_{v_X}^X$ and another containing $P_1^Y, \ldots P_{v_Y}^Y$.

Once a sample space has been identified with a set of sentences, the next step is to translate the model's constraints on probability functions into inductive logical axioms. To show how this can be done I shall now work through an example, considering two statistical models that are prominent in Gelman and Shalizi (2012).

Example

Gelman and Shalizi introduce the following models as exemplars of statistical research in practice at (Gelman and Shalizi, 2012, §2.1). The models were part of an actual investigation into how voting behaviour is related to income in different states of the USA.

$$Pr(y = 1) = logit^{-1}(\mathrm{a}_s + \mathrm{b}x) \qquad \text{(Model One)}$$

and

$$Pr(y = 1) = logit^{-1}(\mathrm{a}_s + \mathrm{b}_s x) \qquad \text{(Model Two)}$$

The voting models in standard presentation

In these models y, s and x stand for categories of random variables $y_1, \ldots, y_n$, $s_1, \ldots, s_n$ and $x_1, \ldots, x_n$, for some very high number n.[1] The indices $1, \ldots, n$ represent voting acts.

Each random variable in the y category has $v_y = 2$ possible values, 1 and 0, representing their corresponding vote going to party 1, the Republicans, or party 0, the Democrats. Variables in category s have $v_s = 50$ values, representing states in which a vote might occur, and those in category x have $v_x = 5$ values

---

[1]In practice, for reasons of mathematical convenience, these sequences of random variables might be taken to be countably infinite.

$1, 2, 3, 4$ and 5 representing the income quintiles to which the person who casts a vote might belong.

The letters a_i are real-valued parameters representing the degree to which, in state i, a vote is more likely to be for the Republicans. The letter b is a real-valued parameter representing a degree to which income-level makes voting Republican more likely. In model two, b is replaced by multiple parameters b_i representing state-specific income effects.

The term $logit^{-1}$ stands for the logistic function, which is a transformation used to make the linear constraints on the right hand sides of the equations in each model apply to the probabilities on the left hand sides.

Both models have the same underlying sample space $\mathcal{S}_{mod}$, which contains all possible configurations of $y_1, \ldots, y_n$, $s_1, \ldots, s_n$ and $x_1, \ldots, x_n$.

The models differ in the sets of probability functions that they allow over $\mathcal{S}_{mod}$, however. These differences are made explicit in the equations above.

The voting models as systems of inductive logic
To construe the two models as systems of inductive logic, we must first find a language L_{mod} corresponding to the sample space $\mathcal{S}_{mod}$.

L_{mod} must have a family of predicates for every category of random variables, each with one member for every possible value of the variables in the corresponding family. In other words, L_{mod} must have $v_y + v_s + v_x = 57$ predicates, organised into a party family R_1^y, R_0^y, a state family $R_1^s, \ldots, R_{50}^s$ and an income quintile family $R_1^x, \ldots, R_5^x$. A basic assumption must stipulate that each constant instantiates exactly one predicate from each family.

To recreate the models' sets of probability functions, we must find axioms characterising two different sets of measure functions $\mathcal{M}_{mod_1}$ and $\mathcal{M}_{mod_2}$. Clearly, probabilism is required to make sure that each set contains only probability functions. Less clearly, constant exchangeability is also needed to make sure that the order of the votes in the sample is ignored. This condition is implicit in the standard presentation's specification of the models in terms of the categories y, s and x rather than the individual random variables $y_1, y_2, \ldots$. Finally, the restrictions stipulated above as equations must be reproduced as the following axioms:

$$m(R_1^y a_i \wedge R_j^s a_i \wedge R_k^x a_i) = logit^{-1}(a_j + b_k) \qquad \text{(Axiom 1)}$$

and

$$m(R_1^y a_i \wedge R_j^s a_i \wedge R_k^x a_i) = logit^{-1}(a_j + b_{k,j}) \qquad \text{(Axiom 2)}$$

Given a suitable specification of the parameters[2], these axioms reproduce the

----

[2]In the axioms, the b parameters must be more complicated—a vector in the first case and a matrix in the second—in order to reproduce the effect of the values of the random variables

standard constraints exactly, though the inductive logical notation forces the presentation to be more complicated. The required systems are then system 1, $(L_{mod}, \mathcal{M}_{mod_1})$, and system 2, $(L_{mod}, \mathcal{M}_{mod_2})$.

This example shows that two models that are typical of current statistical research can not only be thought of as systems of inductive logic in an abstract sense, but have a very similar structure to the systems that Carnap studied.

Choosing between the models

Gelman and Shalizi report that model 2 was preferable to model 1 because, on comparison with a real-life sample recording facts about actual voting events, the parameters in model 1 could not be adjusted so as to achieve a satisfactory fit: the extra flexibility of model 2's extra parameters, representing state-specific income effects, was required. Since the knowledge that enabled the researchers to judge whether a satisfactory fit had been achieved was not represented in any statistical model, Gelman and Shalizi argue that this example vindicates their rejection of **RV**.

Applying their falsificationist account, they claim that model 1 had been compared with data, found not to resemble the data source enough to make future predictions reliable, and therefore rejected. Model 2, in contrast, passed this test and was therefore provisionally accepted.

According to the Carnapian account of statistical model choice, which model is preferable depends on which out of system 1 and system 2 is the most apt formalisation of a collection of inductive assumptions. Since the two models are otherwise identical, the question is which out of axioms 1 and 2 does a better job. Axiom 1 is simpler, as it has fewer parameters. On the other hand, the two axioms represent different assumptions: axiom 1 postulates a single income effect, whereas axiom 2 represents the assumption that the effect may vary between different states. Given the real life sample that Gelman and Shalizi describe, the assumption formalised by axiom 1 is clearly not entirely appropriate. Which system is best will depend on how system 1's comparative disadvantage of representing a less appropriate inductive assumption is weighed against its advantageous simplicity.

Both the falsificationist account and the Carnapian account can claim to be able to reconstruct the reasoning that led real researchers to prefer model 2 to model 1. For the falsificationist account, model 1 was falsified, but model 2 was not. For the Carnapian account, model 1 was found to represent an inapt inductive assumption, whereas model 2 was found to be a useful explication.

---

x_i in the original models being numerical.

6.4 Advantages of the Carnapian account

While the Carnapian account and the falsificationist account can both account for examples like the one above, the Carnapian account has two impotant advantages. It is more technically fruitful than the falsifcationist account, and has a better developed philosophical foundation.

6.4.1 Technical fruitfulness

The Carnapian account is technically fruitful, as considering statistical models' corresponding systems of inductive logic can allow statisticians to express some of their reasons for choosing models more articulately than would be possible otherwise.

When it is possible to find a model's corresponding system of inductive logic, and that system explicates a particular inductive assumption, the model can be defended or attacked according to how appropriate that assumption is. In this way, previously unarticulated justifications can be replaced by transparent ones based on explicitly stated principles.

Finding statistical models' corresponding systems of inductive logic can also help to make clear whether, in a particular case, a model's adoption is primarily based on analytical convenience or if it represents a substantive judgement about the nature of the scientific problem being confronted.

This advantage has previously been identified by Zabell, who argues at in (Zabell, 2011, p. 291) that statisticians' reasons for choosing models often lack such clarity.

To demonstrate this technical fruitfulness, I present below two examples of important classes of statistical models that are characterised by systems of inductive logic that formalise natural inductive assumptions.

Independent, identically distributed random variables and constant exchangeability

De Finetti's theorem identifies a duality between an important class of statistical models, the 'IID models', and systems of inductive logic with unary predicate languages and sets of measure functions characterised by the axioms probabilism and constant exchangeability.

An IID model consists of a sample space $\mathcal{S}$ generated by an infinite sequence of random variables $X_1, X_2, \ldots$, together with a probability function $\mathbf{Pr} : A \to [0, 1]$, known as the 'De Finetti prior', from the set A of independent and identical distributions over $X_1, X_2, \ldots$ to the unit interval. A probability function pr is a member of A if and only if it satisfies the following properties:

Independence
$pr(X_i = v_i, X_j = v_j) = pr(X_i = v_i) \cdot pr(X_j = v_j)$ for any natural numbers i and j and possible values v_i and v_j.

Identical Distribution
$pr(X_i = v_k \mid \theta) = pr(X_j = v_k \mid \theta)$ for any natural numbers i and j, possible value v_k, and any $\theta \in \mathcal{S}$.

De Finetti's representation theorem shows, for unary predicate languages with infinitely many constants, that every IID model has a unique corresponding inductive logical measure function satisfying the axioms of probabilism and constant exchangeability. In the other direction, every probabilistic, exchangeable measure function has a unique corresponding IID model.

See (Paris and Vencovská, 2015, Ch. 9) for a full modern presentation of De Finetti's theorem. See (Carnap, 1980, p.217) and Jeffrey (1971) for discussions by Carnap and his coworker Richard Jeffrey.

Thanks to the correspondence between IID models and the axioms of probability and constant exchangeability, statisticians' choices to use IID models can be justified or criticised according to whether these axioms represent reasonable inductive assumptions in the relevant scientific situation. For example, in circumstances where the order of experiments carries information, constant exchangeability is not a reasonable inductive assumption, giving a reason not to use IID models, according to the Carnapian account.

This recommendation seems to have been adopted in principle by the statistical community. For example, Gelman et al. write:

> The usual starting point of a statistical analysis is the (often tacit) assumption that the n values y_i [in this context these indices are analogous to constants of a logical language] may be regarded as *exchangeable*... A nonexchangeable model would be appropriate if information relevant to the outcome were conveyed in the unit indexes... Generally, it is useful and appropriate to model data from an exchangeable distribution as independent and identically distributed...
> (Gelman et al., 1995, p.6, round parenthesis and italics original, square parenthesis added)

The Carnapian account gives a plausible interpretation of what is meant by the condition that states "may be regarded as exchangeable": it means that constant exchangeability and probabilism explicate appropriate inductive assumptions.

Dirichlet models and Johnson's sufficientness postulate

Every IID model whose De Finetti prior is a member of the Dirichlet family corresponds to an inductive logical measure function that satisfies, in addition

to probabilism and constant exchangeability, the axiom 'Johnson's sufficientness postulate'.

This relationship was demonstrated in principle by Johnson (1932) in the 1930s. (Kemeny, 1963b, § 4) later used Johnson's sufficientness postulate, probabilism and constant exchangeability to characterise the so-called 'continuum of inductive methods', but did not explicitly connect these inductive methods with IID models with Dirichlet distributions. Zabell (1982) makes this connection (see equation 2.14), as well as presenting a more rigorous and general version of Johnson's proof.

The same reasoning that favours using IID models when constant exchangeability is justified extends to this more specific case: statisticians following the Carnapian account can therefore adopt IID models with Dirichlet priors on a principled basis. The scientific situation they face must render appropriate the inductive assumption that the satisfaction ratio of the values of the variables in some samples, together with the sample sizes, should be the only factors relevant to predictions of future values of the variables.

According to (Zabell, 2011, p.292), future work may produce more correspondences of this kind, allowing statisticians to clarify the assumptions underlying other choices of models. Such research can only add to the Carnapian account's technical fruitfulness.

6.4.2 Philosophical foundations

Carnap's philosophy is in remarkable harmony with the actual practice of statistics, as reported by Gelman and Shalizi. The Carnapian account of model choice therefore has a sounder philosophical foundation than the falsificationist account.

This harmony confers two important advantages on the Carnapian account. First, it shows that statisticians can adopt the Carnapian account without committing to unfamiliar or implausible philosophical positions. It therefore constitutes evidence of the Carnapian account's feasibility. Second, it shows that the Carnapian account is part of a more general underlying philosophy of science. This kind of well-connectedness is advantageous as it allows similarities between foundational problems in statistics and other areas to be more easily identified and exploited.

In constrast, the philosophy of Karl Popper, which underpins Gelman and Shalizi's falsificationist account, is in tension with what they say about the practice of statistics. The falsificationist account therefore has none of these advantages.

Carnapian harmony
Gelman and Shalizi report that, in practice, statistical models often have functions other than representing knowledge. Carnap argued that systems of induc-

tive logic need not represent beliefs in order to be useful explicata. Carnap's expression of this point of view is very similar to Gelman and Shalizi's:

> The adoption of an inductive method is neither an expression of belief nor an act of faith, though either or both may come in as motivating factors. An inductive method is rather an instrument for the task of constructing a picture of the world on the basis of observational data. . .
>
> (Carnap, 1950b, §18)

> [A choice of inductive method] will take into consideration . . . the truth-frequency of predictions and the error of estimates; further, the economy in use, measured by the simplicity of the calculations required; maybe also aesthetic features, like the logical elegance of the definitions and rules involved.
>
> (Carnap, 1952a, p.55)

> . . . the model, for a Bayesian, is the combination of the prior distribution and the likelihood, each of which represents some compromise among scientific knowledge, mathematical convenience and computational tractability. . . . we do not have to worry about making our prior distributions match our subjective beliefs. . .
>
> (Gelman and Shalizi, 2012, p.19-20)

Similarly, Carnap would have been at ease with other aspects of the practice of statistical research that Gelman and Shalizi describe. Carnapian inductive logicians can safely take into account "respects that matter to us", tolerate some degree of non-resemblance between predicted and observed data and use informal assessments of whether or not "reasoning based on the model will be reliable" in the future as an important desideratum for choosing between systems of inductive logic.

Carnap thought that inductive logic was essentially a tool for formalising scientific assumptions. From this pragmatic point of view it is only natural that systems of inductive logic should be evaluated in a way that depends on the priorities of the investigator. He would also have seen nothing problematic about failing to reject statistical models that are not completely satisfactory: this is exactly his view of a system of inductive logic that he advocated:

> It will not be claimed that $c^\star$ [the only confirmation function allowed by the system of inductive logic that Carnap was advocating] is a perfectly adequate explicatum. . . For the time being it would be sufficient that $c^\star$ be a better explicatum than the previous methods.
>
> (Carnap, 1950b, p.563, square parentheses added)

Finally, just as, according to Gelman and Shalizi, practising statisticians see no problem in leaving the Bayesian inferential framework in order to evaluate statistical models, Carnap saw no problem with leaving the inferential framework

of inductive logic in order to evaluate systems of inductive logic. Such external questions, he thought, ought to to be answered using testing and the experience of specialists rather than general philosophical proscriptions:

> To decree dogmatic prohibitions of certain linguistic forms instead of testing them by their success or failure in practical use, is worse than futile... Let us grant to those who work in any special field of investigation the freedom to use any form of expression which seems useful to them; the work in the field will sooner or later lead to the elimination of those forms which have no useful function. (Carnap, 1950a, §5)

Popperian tension

In contrast to their natural fit with Carnap's philosophy, the prevailing views of practising statisticians, as reported by Gelman and Shalizi, are hard to square with the philosophy of Karl Popper.

Contrary to Gelman and Shalizi's view that measures of resemblance between models and data sources should take into account the investigation's priorities, Karl Popper argued that disagreement between scientific systems and empirical facts should be investigated in an objective way that does not depend on what matters to scientists. The nature of Popper's view is clear from this passage, where he criticises a proposal by Reichenbach to define statistical hypotheses' probability as the relative frequency with which they have previously been instantiated:

> ... the suggested definition would make the probability of a hypothesis hopelessly subjective: the probability of a hypothesis would depend upon the training and skill of the experimenter rather than upon objectively reproducible and testable results. (Popper, 1959, p.256)

Popper also thought that the conditions under which a theory should be rejected were sharply defined: he claimed at (Popper, 1959, p.56) that rejection should occur whenever a reproducible effect that is inconsistent with the theory is discovered. This stance is at odds with Gelman and Shalizi's portrayal of the prevailing view among practising statisticians, according to which some kinds of non-resemblance between models and data sources should be tolerated.

Gelman and Shalizi seem to acknowledge this divergence from Popper's views on the question of model-checking, writing at (Gelman and Shalizi, 2012, p.28) that "Popper's specific ideas about testing require, at the least, substantial modification". However, there are further points of tension between Popper's position and the statistical mainstream.

Popper thought that the main concern of proper scientific research should be attempting to demonstrate that theories are false. He saw activities that do not assist this process as not strictly scientific. It seems difficult, following such

an approach, not to construe the adoption and rejection of statistical models as expressions of belief and disbelief, or to find conscionable the view, which Gelman and Shalizi say is widespread among statisticians, that desiderata with little to do with truth and falsity, such as convenience or tradition, should play important roles in statistical research.

Finally, Popper was an anti-inductivist: he thought that the adoption of scientific theories should never depend on judgements about the accuracy of their future predictions based on past observations. Popper believed that such inductive judgements were unscientific:

> Induction, i.e. inference based on many observations, is a myth. It is neither a psychological fact, nor a fact of ordinary life, nor one of scientific procedure. (Popper, 1962, p.53)

In contrast to this view, Gelman and Shalizi report that statisticians do and should make use of judgements, based on past observations, about the future reliability of reasoning based on statistical models. Such judgements, they note, are required in order to determine how to measure resemblance between models and data sources and how much non-resemblance to tolerate before rejection. According to Gelman and Shalizi's assessment, then, practising statistician seem to be inductivists in Popper's sense.

For all of these reasons, the practice of statistics seems to be more than superficially incompatible with Popper's philosophy of science. Practising statisticians seem to disagree with Popper on the fundamental questions of why theories should be accepted and rejected and whether science should be inductive.

6.4.3   Conclusion

In summary, the Carnapian account of statistical model choice seems superior to the falsificationist account found in Gelman and Shalizi (2012). It is technically fruitful, potentially allowing more articulate expression of the reasons behind the selection of particular models. Clearly stated modelling assumptions can only improve statistical research. In addition, unlike the falsificationist account, the Carnapian account has a viable underlying philosophy that is in tune with Gelman and Shalizi's claims about the practice of statistics.

It is difficult to think of another philosophical research programme that can play the same role. While subjective Bayesian and objective Bayesian inductive logic may have interesting applications to philosophical problems in statistics as well, neither is flexible enough to generate a general account of statistical model choice. Subjective Bayesian inductive logic is committed to the conditional probability method of representing evidence, and so cannot countenance the kind of informal evidential reasoning that Gelman and Shalizi claim is crucial in real-world statistical model-choice. This kind of reasoning is readily accommodated by Carnapian inductive logic, which allows evidence to be represented

informally in the choice of systems of inductive logic. Statistical investigations are also a poor match with objective Bayesian inductive logic because they often make use of the conditional probability method for representing possible ways of reacting to data. Objective Bayesian inductive logic, unlike Carnapian inductive logic, rejects this method of representing evidential reasoning.

These advantages show that Carnapian inductive logic can play a useful role in the contemporary debate over model-choice in statistics. This is a concrete example of why substantively reappraising it would be a good idea.

It is interesting to note that Carnap himself saw one of the principal application of his research programme as being to statistics. He argues in (Carnap, 1950b, § 49A), entitled 'Theoretical Usefulness of Inductive Logic in Science' that Carnapian inductive logicians should aim to formalise the inductive assumptions that arise in statistical investigations. He described the benefits of doing so as follows:

> It may be expected that mathematical statistics will thereby gain for the first time a solid foundation, a systematic unity of its various methods and a clarity and exactness of its basic concepts. . . . In spite of the great wealth in methods and results achieved in modern mathematical statistics . . . it is clearly in need of the theoretical virtues just mentioned. . . (Carnap, 1950b, p. 244)

The usefulness of the Carnapian account even in modern discussions about statistics goes some way towards vindicating Carnap's assessment.

Chapter 7

Conclusion

This thesis has argued that a reappraisal of Carnapian inductive logic would be a good idea, and begun to carry one out. In this section I summarise the preceding discussions and outline a research agenda for future reappraisers.

7.1 The current state of the reappraisal

This thesis aimed to begin a historical reappraisal of Carnapian inductive logic by putting forward a new account of its nature and development, and also to start a substantive reappraisal by pointing out ways in which Carnapian inductive logic remains philosophically interesting.

7.1.1 Historical reappraisal

My novel historical contributions came mostly in chapters 2 and 3.

Chapter 2 presented and defended a novel reading of Carnapian inductive logic according to which it is characterised by commitments to explication, systems of inductive logic and tolerance. This account differs sharply from previous ones, which have tended to characterise Carnapian inductive logic as a highly intolerant research programme which tried to identify one uniquely 'true' or 'logical' system of inductive logic.

Chapter 2 also documented the ways in which Carnapian inductive logic changed as it developed, arguing that these changes were less drastic than is currently thought. Whereas much of the secondary literature emphasises purported differences between early and late Carnapian inductive logic, I argued that the programme is noteworthy for the degree to which it remained consistent with its initial aims and methods.

Chapter 3 defended Carnapian inductive logic, understood according to my reading, against what I judeged to be the most important of its many critiques. It is widely believed that the strength of these critiques caused Carnapian inductive logic to become "moribund" (Williamson, 2002, p.210), exposed a "fatal flaw"(Hawthorne, 2014, §), showed that it had "faltered"(Sarkar, 2006, p.87) and so on. My argument suggests that this was not the case.

My thesis made another novel historical contribution in section 5.4.4 on the assimilative approach. There I argued, contrary to the conventional historical account of Carnapian inductive logic, that Carnap considered an alternative to the conditional-probability method for representing evidence, citing a previously under-appreciated draft essay.

7.1.2   Substantive reappraisal

This thesis began a substantive reappraisal of Carnapian inductive logic by arguing that Carnapian inductive logic is a viable and interesting research programme.

Chapter 2 argued that Carnapian inductive logic was successful according to Lakatos's methodology of scientific research programmes. This amounted to a prima facie argument for its viability, but, I noted, should not be relied on too heavily. For one thing, Lakatos's methodology is not without its critics. For another, three important issues gave independent reasons to doubt whether Carnapian inductive logic remained relevant. First, it has been heavily and apparently successfully critiqued. Second, one might doubt whether Carnapian inductive logic can hope to survive in the modern conceptual ecosystem, given the proliferation of similar research programmes that might outcompete it. Third, it is not obvious that Carnapian inductive logic can solve any contemporary problems. The subsequent chapters addressed these issues.

Chapter 3 argued that the most important of the many critiques of Carnapian inductive logic can be overcome. Lakatos set out explicitly a mistaken assumption that seems to underlie many of the critiques, according to which Carnapian inductive logic aimed to discover certain a priori, theoretically justified truths. This kind of critique can be dispelled through a careful reading of Carnap's writing about inductive logic, which shows clearly that theoretical justification was not its aim. The 'grue' problem can be addressed by noting that Carnapian inductive logic can formalise reasoning about both projectable and non-projectable properties, together with further close reading to show that Carnapian inductive logic did not aim to produce interpretation-neutral axioms. Worries stemming from concerns about Carnapian inductive logic's treatment of exchangeability, partial entailment, reliability in the limit and universal hypotheses can also be assuaged.

Chapters 4 and 5 addressed the question of Carnapian inductive logic's relationship with rival research programmes. We saw that one rival programme, sub-

jective Bayesian inductive logic, makes philosophically dubious commitments to probabilistic necessity, the claim that all rational states of partial belief can be represented by probability functions, and probabilistic sufficiency, the claim that every state of partial belief that can be so represented is rational. Carnapian inductive logic is preferable to subjective Bayesian inductive logic from a philosophical point of view as it avoids making such controversial claims about rationality. On a methodological level, Carnapian inductive logic can formalise all of the inductive assumptions that subjective Bayesian inductive logic can, but the reverse is not true. Subjective Bayesian inductive logic's commitment to representing evidence using the conditional probability method prevents it from being as comprehensive as Carnapian inductive logic. While subjective Bayesian inductive logic is arguably simpler than Carnapian inductive logic, on balance Carnapian inductive logic seems methodologically preferable.

Objective Bayesian inductive logic has advantages and disadvantages compared to Carnapian inductive logic. One disadvantage is that, like subjective Bayesian inductive logic, it is committed to the philosophically dubious claim of probabilistic necessity. Objective Bayesian inductive logic's other commitments make it difficult for it to represent certain kinds of evidence, in particular evidence according to which some events are relevant to others. On the other hand objective Bayesian inductive logic is more specific than Carnapian inductive logic, leaving less room for the inductive logician's discretion: this is a methodological advantage in certain cases. I concluded that whether objective Bayesian or Carnapian inductive logic is preferable will depend on the nature of the problems to which they are applied.

I also argued in chapter 5 that, as well as being used to define an alternative form of inductive logic, the norms of objective Bayesian epistemology can also be formalised within Carnapian inductive logic. This 'assimilative approach' to relating Carnapian inductive logic and objective Bayesian epistemology could usefully inform philosophical debates about objective Bayesian epistemology by clarifying its commitments and allowing the introduction of pragmatic arguments. However, in order to formalise the norms adequately, Carnapian inductive logic must be extended to rich domains whose members can represent both states of the world and states of evidence that take the form of sets of probability functions.

Chapter 6 presented an application of Carnapian inductive logic to the problem of philosophically describing statistical model-choice. I argued in favour of the Carnapian account of statistical model choice, which draws an analogy between this activity and that of choosing between systems of inductive logic. This account is consistent with the practice of statistical research, as reported in a widely-cited paper, and is preferable to an alternative falsificationist account as it is more technically fruitful and based on a more appealing underlying philosophy.

Taken together, the arguments in these chapters provide a sound basis for a substantive reappraisal of Carnapian inductive logic. Not only was it progressive

according to Lakatos, it is also able to overcome its critiques, compete with its rivals and inform contemporary philosophical debates about objective Bayesian epistemology and statistical model-choice.

7.2   The case for ongoing reappraisal

The tasks of motivating and practising reappraisal are, I believe, complementary. The more successful a new programme of reappraisal is, the more attractive a prospect participating in it becomes. To a large extent, then, my argument that Carnapian inductive logic should be reappraised is encapsulated by the reappraisal that I have attempted to begin.

However, it remains conceiveable that a reader, even though they are convinced by all of the preceding arguments, might still be hestitant about committing their time and attention to the reappraisal of Carnapian inductive logic. They might be worried that there is not very much more to say about it. In order to allay this kind of doubt, I outline below a research agenda following on from the arguments in this thesis.

7.2.1   Reassessing the context of Carnapian inductive logic

In order to reappraise a philosophical research programme properly, its historical and intellectual context must be thoroughly understood. There have been great advances in our understanding of Carnap and the logical positivists over the last few decades, but there is still much to be gained from further research. Here are some examples of what I think are interesting and open historical questions about Carnapian inductive logic.

First, it is interesting that, after Carnap's death in 1970, few of his co-workers carried on working on Carnapian inductive logic. For example, John Kemeny, Yehoushua Bar-Hillel and Haim Gaifman pursued research into largely unrelated topics; Richard Jeffrey developed a different philosophical programme that was more inclined towards subjective Bayesian epistemology.

I argued in chapter 3 that Carnapian inductive logic survived its many critiques. Given that this was the case, it is an interesting and unsolved question why Carnap's coworkers discontinued his research programme. An alternative explanation, besides the revelation by critics of its fundamental unviability, is required.

Secondly, it would be informative to know more about how Carnap's views on inductive logic relate to those of the other logical positivists. Hans Reichenbach, Herbert Feigl and Carl Hempel corresponded with Carnap frequently during the time when he was working on Carnapian inductive logic, and all had well-documented views about the philosophical problem of induction. See Salmon

(1991) for an account of Reichenbach and Feigl's views, and Hempel (1981) for those of Hempel. Nonetheless, it is unclear exactly what they thought about Carnapian inductive logic, and how their views developed in dialogue with Carnap. Although Otto Neurath died before Carnap published on inductive logic, it would also be very interesting to know the extent to which he and Carnap exchanged ideas about probability and induction - see Lehrer (1993) for some interesting informed speculation on this point.

Finally, it would be helpful if Carnap's early views on induction, before the 1940s, were better understood.

It will surely become easier to answer these questions in the near future, as various investigations into Carnap's early views are ongoing. A growing amount of archival material is available online at the university of Pittsburgh library website. Electronic copies of other material are available on request. In addition, several historical research projects on Carnap are ongoing. The project 'Early Carnap in Context - Three Case Studies and the Diaries' is seeking to translate diaries that Carnap kept between 1910 and 1935 and assess their significance. Another research project is seeking to collate Carnap's complete works, with commentaries by experts.

7.2.2 Carnap's conception of logic

As we have seen, Carnap was keen to stress parallels between his research programme and deductive logic. I noted that these parallels make sense with respect to Carnap's conception of logic, according to which systems of inductive logic qualify as logical on the basis that they are systems of formal rules specifying what follows from what within a language. However, I did not dwell at length on Carnap's conception of logic, either to track its historical development or to assess its viability. In my opinion, the question whether or not Carnapian inductive logic merits reappraisal does not depend on whether it is correct to call it a 'logic'. Nonetheless, the nature and viability of Carnap's conception of logic remain pertinent topics for a reappraisal of Carnapian inductive logic.

On a historical level, it would be interesting to find out how Carnap's mature conception of logic, according to which systems of inductive logic qualify as logical, relates to his earlier views. As far as I am aware, no work connecting Carnap's views about inductive logic with the development of his views about logic in general has yet been undertaken. At a conceptual level, a natural question to ask is whether Carnap's conception of logic is viable, or at least salvageable. Carnap's views on logic have been critically appraised in various recent works, such as Peregrin (2015), Awodey (2015) and Goldfarb and Ricketts (1992). A special issue of *Synthese* entitled 'Carnap on logic' is forthcoming. Again, to my knowledge, these works do not engage with Carnap's views on inductive logic. Clearly there is much fertile ground for further work in this area.

In addition to questions about Carnap's conception of logic, one might also ask whether Carnapian inductive logic would qualify as logical according to alternative conceptions. Answering this question may help to explain the frequency with which Carnapian inductive logic has been attacked on the grounds that it is not genuinely logical.

7.2.3   Other ways of evaluating research programmes

Chapter 2 applied Lakatos's methodology of scientific research programmes, arguing that, according to this methodology, Carnapian inductive logic was a progressive research programme. However, as I noted in that chapter's conclusion, there are several outstanding objections to Lakatos's methodology. For the purposes of this thesis, the status of Lakatos's methodology was not of primary importance, as its role was only to provide prima facie, rather than conclusive, evidence that Carnapian inductive logic is worthy of substantive reappraisal.

However, a more comprehensive evaluation would be very interesting. In order to come to a final conclusion about Carnapian inductive logic's success as a research programme, an evaluative framework will need to be employed that either addresses the objections to Lakatos's methodology or else departs from it.

The version of Lakatos's methodology outlined in (Corfield, 2003, Ch. 8) is a good candidate for the first option. Corfield modifies Lakatos's approach by changing its focus from truth-directed propositions to vaguer 'ideas'. If the methodology of scientific research programmes, so modified, survives the critiques that have been aimed at Lakatos's version, then it could usefully be used to evaluate Carnapian inductive logic. Alternatively, an evaluative framework based on Laudan's work on 'research traditions', as sketched in Laudan (1978), or the 'philosophical grammar' approach of Chang (2011) might be preferable.

7.2.4   Carrying out the assimilative approach

The assimilative approach to relating objective Bayesian epistemology and Carnapian inductive logic was incomplete, as I did not fully explain what would be the abstract form of an inductive logical domain that can represent both states of the world and states of evidence. There are many technical questions to answer, such as what kind of algebraic structure this kind of domain should have. If these technical challenges can be overcome, then it will be possible to compare axioms representing aspects of objective Bayesian epistemology directly with axioms representing alternative epistemological views. This will open the possibility for new applications of Carnapian inductive logic to philosophy in the form of novel arguments for or against objective Bayesian epistemology, and perhaps even applications to science in the form of useful new systems of inductive logic.

7.2.5 Imprecise probability

Carnapian inductive logic has a potential applications to the problem of ascertaining the status of imprecise approaches to representing uncertain reasoning. The question of whether, or in what circumstances, imprecise approaches are legitimate has become a vexed problem in recent formal epistemology. Construing the question as one of applied Carnapian inductive logic, I claim, might clarify the situation.

An imprecise model of uncertain belief represents a state of uncertain belief using a set of measure functions—a 'representor' in the terminology of Van Fraasen (1990)—rather than a single such function.

This approach has been attacked and defended on both methodological and epistemological grounds.

Epistemological debate
The epistemological debate focuses on sharpness, the claim that, if the objects of a rational state of partial belief are represented by a domain D, then the state of belief itself should be representable by a single real-valued measure function $m : D \to \mathbb{R}$.

Proponents of imprecise models argue that sharpness is not a genuine rationality condition because it incorrectly identifies certain rational states of partial belief as irrational. In particular, they claim, rational states of partial belief may lead to dispositions to bet for and against certain propositions at different rates. Walley (1991) argues at length that such states of belief are, in general, best represented by sets of probability functions. In addition it has been argued, for example in Keynes (1921), that rational degrees of belief can be incomparable in principle. Sharpness forces that all degrees of belief to be comparable and, according to this argument, consequently mis-classifies such states of belief as irrational.

Epistemological opponents of imprecise models argue for sharpness by trying to demonstrate that non-sharp states of partial belief have undesirable features. For example, Elga (2010) uses a diachronic betting scenario to argue that agents with non-sharp states of belief tend to behave unwisely. Bradley and Steele (2014) disputes this argument.

Methodological debate
Defenders of imprecise models point to distinct advantages of representing states of belief using sets of probability functions. Dubois et al. (1996) argues that imprecise models are required in order to represent higher-level uncertainty; (Walley, 1991, § 1.1.4) claims that imprecise models can be used to produce inferences that are more robust than would be possible otherwise.

Opponents, on the other hand, argue that imprecise models are difficult to connect with formal models of decision-making. Whereas the question of how agents with sharp states of partial belief should make decisions is comparatively well-understood theoretically, the theory of how imprecise agents should make decisions is still fairly young.

How Carnapian inductive logic can help

Both the epistemological and methodological debates over imprecise approaches to representing uncertain reasoning seem to lack an over-arching framework for deciding whether each kind of argument is decisive. This, I claim, is what Carnapian inductive logic may be able to provide.

Imprecise models and precise models of partial beliefs can both be thought of as systems of inductive logic. Both sorts of model come with the right kind of domain, and ready made sets of measure functions. An imprecise model consisting of a domain of propositions D together with a representor P can straightforwardly be construed as the system of inductive logic (D, P). A precise model with domain D and representing measure function m can naturally be construed as the system of inductive logic $(D, \{m\})$

Since both precise and imprecise models of uncertain reasoning can be thought of as systems of inductive logic, it is possible to apply Carnap's standards of evaluation. It therefore seems feasible to formulate a Carnapian account of choosing models of partial belief that encompasses both precise and imprecise models. According to this account models would be judged more or less legitimate to the extent that their corresponding systems of inductive logic are good ones, according to Carnapian standards. This would provide a framework within which debates about the epistemological and methodological pros and cons of each kind of model could proceed.

Using Carnapian inductive logic in this way, rather than merely comparing imprecise and precise approaches as they appear naturally, would allow a direct, systematic and unambiguous comparison of the two approaches to formalising uncertain reasoning.

As in the case of statistical model-choice, other forms of inductive logic do not seem to be capable of playing the same role. Subjective and objective Bayesian inductive logic are committed to probabilistic necessity, according to which rational states of partial belief are probabilistic and consequently sharp. These two frameworks therefore presume that imprecise models of rational partial belief are epistemologically illegitimate, and so cannot help with the epistemological debate. These forms of inductive logic also have little to contribute to the debate over imprecise models' methodological legitimacy because they are not capable of formulating such models.

7.2.6 Double-counting

Another potential application of Carnapian inductive logic is to an ongoing debate over 'double-counting' in the context of climate models that is discussed in Steele and Werndl (2013). Double-counting is the practice of using the same data to both calibrate parameters within a climate model and to evaluate it against alternative models. For example, Steele and Werndl refer to a study in which historic data about the Earth's temperature was used to find parameters that represent, within two models, the Earth's overall climate sensitivity and the cooling effect caused by aerosols. The same data was then used to evaluate the two models: one model was judged to be better than the other on the grounds that, after this process of calibration, its instances achieved a better fit. Some philosophers and climate scientists hold that using data twice in this way should be avoided where possible, whereas others, including Steele and Werndl, argue that there is nothing intrinsically wrong with it.

A similar situation can occur within Carnapian inductive logic. On the one hand, evidence about experimental results can be used to narrow down the choice of values of a parameter within a parameterised system of inductive logic. This practice is described at (Carnap, 1952a, p. 54), where Carnap advocates excluding certain values of the parameter λ in response to evidence about experimental results. This kind of procedure is analogous to the calibration of climate models using data. On the other hand, Carnap also thought that experimental evidence can be used to evaluate inductive logical axioms: at (Carnap, 1971a, p. 120) he recommends using the axiom of constant exchangeability in cases where "experience shows. . . statistical independence." Carnap seems to have been comfortable with the inductive logical equivalent of double-counting, as the use of experimental evidence to choose a value of λ in Carnap (1952a) would have to be made after an experience-based choice to use the axiom of constant exchangeability, which is taken for granted in this work.

Paying close attention to Carnap's discussion of these inductive logical roles for experimental evidence might make it possible to generalise of Steele and Werndl's argument in favour of double-counting in climate science. Steele and Werndl's argument assumes that model-choice takes place within a Bayesian framework: all factors contributing to model-choice are taken to be represented in a higher-level rational-belief model. These assumptions could potentially be weakened by reconstructing the argument within Carnapian inductive logic. This would be desirable as the Bayesian framework is not universally accepted among climate scientists.

7.3 Conclusion

Carnapian inductive logic has much to offer contemporary philosophy. It is conceptually sound, historically interesting and has as-yet unexplored connections

with many important problems. In these circumstances, there only seems to be one sensible course of action: let's reappraise Carnapian inductive logic!

Bibliography

Awodey, S. (2015). Carnap and the invariance of logical truth. *Synthese.* http://dx.doi.org/10.1007/s11229-015-0781-6.

Awodey, S. and Carus, A. (2001). How Carnap could have replied to Gödel. Technical Report CMU-PHIL-123, Carnegie Mellon University.

Bar-Hillel, Y. and Carnap, R. (1953). Semantic information. *British Journal for the Philosophy of Science*, 4(14):147–157.

Blackwell, D. and MacQueen, J. B. (1973). Ferguson distributions via Polya urn schemes. *The Annals of Statistics*, 1(2):353–355.

Boyd, S. and Vandenberghe, L. (2009). *Convex Optimization.* Cambridge University Press, 7th edition.

Bradley, S. and Steele, K. (2014). Should subjective probabilities be sharp? *Episteme*, 11(3):277–289.

Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review*, 122(4):527–575.

Carnap, R. (1928). *The Logical Structure of the World.* University of California Press (quotes from English translation by Carnap, published 1967).

Carnap, R. (1937). *Logical Syntax of Language.* Routledge.

Carnap, R. (1945a). On inductive logic. *Philosophy of Science*, 12(2):72–97.

Carnap, R. (1945b). The two concepts of probability. *Philosophy and Phenomenological Research*, pages 513–532.

Carnap, R. (1946). Remarks on induction and truth. *Philosophy and Phenomenological Research*, pages 590–602.

Carnap, R. (1947a). *Meaning and Necessity: a study in semantics and modal logic.* University of Chicago Press.

Carnap, R. (1947b). On the application of inductive logic. *Philosophy and Phenomenological Research*, 8(1):133–148.

Carnap, R. (1947c). Probability as a guide in life. *Journal of Philosophy*, 44(6):141–148.

Carnap, R. (1948). Reply to Nelson Goodman. *Philosophy and Phenomenological Research*, 8(3):461–462.

Carnap, R. (1950a). Empiricism, semantics, and ontology. *Revue Internationale De Philosophie*, 4(2):20–40.

Carnap, R. (1950b). *Logical Foundations of Probability*. University of Chicago Press. 2nd Edition: published 1962.

Carnap, R. (1951). The problem of relations in inductive logic. *Philosophical Studies*, 2:75–80.

Carnap, R. (1952a). *The Continuum of Inductive Methods*. University of Chicago Press.

Carnap, R. (1952b). Meaning postulates. *Philosophical Studies*, 3(5):65–73.

Carnap, R. (1953). Inductive logic and science. *Proceedings of the American Academy of Arts and Sciences*, 80(3):189–197.

Carnap, R. (1963a). Arthur W. Burks on the philosophy of induction. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, chapter 28, pages 980–983. Open Court publishing co.

Carnap, R. (1963b). An axiom system for inductive logic. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, chapter 26, pages 972–978. Open Court publishing co.

Carnap, R. (1963c). Hilary Putnam on degree of confirmation and inductive logic. In Schilpp, P., editor, *The philosophy of Rudolf Carnap*, chapter 24, pages 983–989. Open Court publishing co.

Carnap, R. (1963d). Intellectual autobiography. In Schilpp, P., editor, *The Library of Living Philosophers - Rudolf Carnap*, pages 3–81. Open Court.

Carnap, R. (1963e). Popper on probability and induction. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, chapter 31, pages 995–998. Open Court publishing co.

Carnap, R. (1966a). The aim of inductive logic. *Studies in Logic and the Foundations of Mathematics*, 44:303–318.

Carnap, R. (1966b). *Philosophical Foundations of Physics*. Basic Books inc.

Carnap, R. (1967a). Letter to Lakatos dated June 29 1967. http://digital.library.pitt.edu/u/ulsmanuscripts/pdf/31735062224211.pdf.

Carnap, R. (1967b). Uncertain Evidence. Ms. "Uncertain Evidence," Gordon Matthews, 1966-1968 Box 107b, Folder 10; Rudolf Carnap Papers, 1905-1970, ASP.1974.01, Special Collections Department, University of Pittsburgh See http://digital.library.pitt.edu/u/ulsmanuscripts/pdf/31735061812735.pdf.

Carnap, R. (1968). Inductive logic and inductive intuition. *Studies in Logic and the Foundations of Mathematics*, 51:258–314.

Carnap, R. (1971a). *A basic system of inductive logic, Part 1*, volume 1 of *Studies in Inductive Logic and Probability*, chapter 2, pages 33–167. University of California Press.

Carnap, R. (1971b). Inductive logic and rational decisions. In Carnap, R. and Jeffrey, R. C., editors, *Studies in Inductive Logic and Probability*, volume 1, pages 7–31. University of California Press.

Carnap, R. (1980). *A basic system of inductive logic, Part 2*, volume 2 of *Studies in Inductive Logic and Probability*, chapter 6, pages 7–156. University of California Press.

Carnap, R. and Shimony, A. (1977). *Two Essays on Entropy*. University of California Press.

Carnap, R. and Stegmüller, W. (1959). *Inductive logik und Warscheinlicheit*. Wien Springer Werlag.

Carus, A. (2010). *Carnap and Twentieth Century Thought*. Cambridge University Press.

Carus, A. (2014). Carnapian rationality. *Synthese*.

Chandler, J. (2014). Subjective probabilities need not be sharp. *Erkenntnis*, pages 1–14.

Chang, H. (2011). The philosophical grammar of scientific practice. *International Studies in the Philosophy of Science*, 25(3):205–221.

Chomsky, N., Dobereiner, D., Hess, J., Richardson, D., and Woodhull, T. (1988). Anarchism (transcript of interview with Noam Chomsky. In Otero, C., editor, *Language and Politics*, pages 166–196. Black Rose books.

Colyvan, M. (2004). The philosophical significance of Cox's theorem. *International Journal of Approximate Reasoning*, 37(1):71–85.

Corfield, D. (2003). *Towards a Philosophy of Real Mathematics*. Cambridge University Press.

Cox, R. (1947). Probability, frequency and reasonable expectation. *American Journal of Physics*, 14(1):1–13.

Creath, R. (1991). Every dogma has its day. *Erkenntnis*, 35:347–389.

De Finetti, B. (1931). Sul significato soggettivo della probabilità. *Fundamenta Mathematicae*, (17):298–329.

De Finetti, B. (1990). *Theory of Probability, a critical introductory treatment*, volume 1. Wiley and Sons.

Dubois, D., Prade, H., and Smets, P. (1996). Representing partial ignorance. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 26(3).

Earman, J. (1989). Concepts of projectability and the problems of induction. *Minnesota Studies in the Philosophy of Science*, 12:220–233.

Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10(5).

Feyerabend, P. (1981). *Problems of Empiricism*, volume 2 of *Philosophical Papers*. Cambridge University Press.

Fitelson, B. (2005). Inductive logic. In Sarkar, S. and Pfeifer, J., editors, *The Philosophy of Science: An Encyclopedia*, volume 1. Routledge.

Friedman, M. (1999). *Reconsidering Logical Positivism*. Cambridge University Press.

Gaifman, H. (1971). Applications of de Finetti's theorem to inductive logic. In Carnap, R. and Jeffrey, R., editors, *Studies in Inductive Logic and Probability*, volume 1, pages 235–236. University of California Press.

Galison, P. (1990). Aufbau/Bauhaus: Logical positivism and architechtural modernism. *Critical Inquiry*.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian Data Analysis*. Chapman and Halll.

Gelman, A. and Shalizi, C. R. (2012). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66(1):8–38.

Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4):650–669.

Glaister, S. (2007). *Inductive Logic*, pages 563–581. Blackwell Publishing Ltd.

Goldfarb, W. and Ricketts, T. (1992). Carnap and the philosophy of mathematics. In *Science and Subjectivity*, The Vienna Circle and 20th Century Philosophy. Akademie Verlag.

Goodman, N. (1946). A query on confirmation. *The Journal of Philosophy*, 43(14):383–385.

Goodman, N. (1947). On infirmities of confirmation-theory. *Philosophy and Phenomenological Research*, 8(1):149–151.

Goodman, N. (1954). *Fact, Fiction and Forecast*. Harvard University Press.

Greaves, H. (2013). Epistemic decision theory. *Mind*, 122(48):915–952.

Hacking, I. (1967). Slightly more realistic personal probability. *Philosophy of Science*, 34(4):311–325.

Hájek, A. (2008). Arguments for - or agains - probabilism? *British Journal for the Philosophy of Science*, 59:793–819.

Hájek, A. (2012). Interpretations of probability. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Summer 2012 edition.

Halpern, J. (1999a). A counterexample to theorems of Cox and Fine. *Journal of Artificial Intelligence Research*, 10:67–85.

Halpern, J. (1999b). Cox's theorem revisited. *Journal of Artificial Intelligence Research*, 11:429–435.

Harman, G. (1983). Problems with probabilstic semantics. In Orenstein, A. and Stern, R., editors, *Developments in Semantics*, pages 237–243. Haven.

Hawthorne, J. (2009). The Lockean thesis and the logic of belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 49–74. Springer Netherlands.

Hawthorne, J. (2014). Inductive logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2014 edition.

Hempel, C. G. (1945). Geometry and empirical science. *The American Mathematical Monthly*, 52(1):7–17.

Hempel, C. G. (1981). Turns in the evolution of the problem of induction. *Synthese*, 46(3):389–404.

Hill, A. (2013). *Reasoning by Analogy in Inductive Logic*. PhD thesis, University of Manchester.

Howson, C. (1997a). Logic and probability. *The British Journal for the Philosophy of Science*, 48(4):517–531.

Howson, C. (1997b). A logic of induction. *Philosophy of Science*, 64(2):268–290.

Howson, C. (2000). *Hume's Problem*. Clarendon Press, Oxford.

Howson, C. (2001). The logic of Bayesian probability. *Applied Logic Series*, 24:137–159.

Howson, C. (2003). Probability and logic. *Journal of Applied Logic*, 1(3):151–165.

Howson, C. (2009). Can logic be combined with probability? *Journal of Applied Logic*, 7:177–187.

Howson, C. (2012). Modelling uncertain inference. *Synthese*, 186(2):475–492.

Howson, C. and Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*. Open Court (La Salle, Ill.).

Jaynes, E. T. (1995). *Probability Theory: the Logic of Science*. Cambridge University Press.

Jeffrey, R. (1975). Carnap's empiricism. *Minnesota Studies in the Philosophy of Science*, 6:37–49.

Jeffrey, R. C. (1966). Goodman's query. *The Journal of Philosophy*, 63(11):281–288.

Jeffrey, R. C. (1967). Logical probability (jeffrey on salmon on carnap). Salmon, Wesley: APA Paper "Carnap's Inductive Logic" and Comments by Jeffrey and Kemeny, 1967 Box 5, Folder 14 Richard C. Jeffrey Papers, 1934-2002, ASP.2003.02, Archives of Scientific Philosophy, Special Collections Department, University of Pittsburgh.

Jeffrey, R. C. (1971). Probability measures and integrals. In Carnap, R. and Jeffrey, R. C., editors, *Studies in Inductive logic and probability*, volume 1. University of California Press.

Jeffrey, R. C. (1973). Carnap's inductive logic. *Synthese*, 25(3-4):299–306.

Jeffrey, R. C. (1988). De Finetti's radical probabilism. Box 15, Folder 15 Richard C. Jeffrey Papers, 1934-2002, ASP.2003.02, Archives of Scientific Philosophy, Special Collections Department, University of Pittsburgh.

Johnson, W. E. (1932). Probability: The deductive and inductive problems. *Mind*, 41(164):409–423.

Joyce, J. (1998). A non-pragmatic vindication of probabilism. *Philosophy of Science*, 65:575–603.

Joyce, J. (2009). Accuracy and coherence: Prosepects for an alethic epistemology of partial belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*. Synthese Library.

Joyce, J. (2010). A defense of imprecise credences in inference and decision making. *Philosophical perspectives*, 24(1):281–323.

Joyce, J. (2011). The development of subjective Bayesianism. In Gabbay, D. M., Hartmann, S., and Woods, J., editors, *Inductive Logic*, volume 10 of *Handbook of the History of Logic*, pages 415–475. North Holland.

Kelly, K., Juhl, C., and Glymour, C. (1994). Reliability, realism and relativism. In Clark, P. and Hale, B., editors, *Reading Putnam*. Blackwell Publishing Ltd.

Kemeny, J. (1955). Fair bets and inductive probabilities. *Journal of Symbolic Logic*, 20(3).

Kemeny, J. (1963a). Analyticity versus fuzziness. *Synthese*, 15:57–80.

Kemeny, J. (1963b). Carnap's Theory of Probability and Induction. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, pages 711–736. Open Court.

Kemeny, J. (1967). Comments on Salmon's paper: "Carnap's inductive logic". Salmon, Wesley: APA Paper "Carnap's Inductive Logic" and Comments by Jeffrey and Kemeny, 1967 Box 5, Folder 14 Richard C. Jeffrey Papers, 1934-2002, ASP.2003.02, Archives of Scientific Philosophy, Special Collections Department, University of Pittsburgh.

Keynes, J. M. (1921). *A Treatise on Probability*. Macmillan and co.

Lakatos, I. (1966). Letter to Carnap. Konstanz archive, Reference RC 027-48-19.

Lakatos, I. (1967a). The criticism of Carnap's research programme. LSE Lakatos archive, reference LAKATOS/13/142.

Lakatos, I. (1967b). Letter to Carnap. Konstanz archive, Reference RC 088-48-03.

Lakatos, I. (1967c). Unpublished manuscript. Konstanz archive, References RC 088-48-10, RC 088-48-11.

Lakatos, I. (1968a). Changes in the problem of inductive logic. In *The Problem of Inductive Logic*. North Holland.

Lakatos, I. (1968b). Criticism and the methodology of scientific research programmes. *Proceedings of the Aristotelian Society*, 69:149–186.

Lakatos, I., editor (1968c). *The Problem of Inductive Logic*. North Holland Publishing Company.

Lakatos, I. (1970). History of science and its rational reconstructions. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pages 91–136.

Lakatos, I. (1973). Popper on demarcation and induction. In Worrall, J. and Currie, G., editors, *The Methodology of Scientific Research Programmes*, volume 1 of *Philosophical papers*. Cambridge University Press.

Lakatos, I. (1978a). *Falsification and the methodology of scientific research programmes*, chapter 1, pages 8–93. Cambridge University Press.

Lakatos, I. (1978b). *Mathematics, Science and Epistemology*, volume 2. Cambridge University Press.

Lakatos, I. (1978c). Science and pseudo-science. In Worrall, J. and Currie, G., editors, *Philosophical Papers*, pages 1–8. Cambridge University Press.

Landes, J. (2009). *The Principle of Spectrum Exchangeability within Inductive Logic*. PhD thesis, University of Manchester.

Landes, J. (2015). Probabilism, entropies and strictly proper scoring rules. *International Journal of Approximate Reasoning*, 63:1–21.

Landes, J., Paris, J., and Vencovska, A. (2008). Some aspects of polyadic inductive logic. *Studia Logica*, 90:3–16. 10.1007/s11225-008-9140-7.

Landes, J. and Williamson, J. (2013). Objective Bayesianism and the maximum entropy principle. *Entropy*, 15(9):3528–3591.

Landes, J. and Williamson, J. (2015). Justifying objective Bayesianism on predicate languages. *Entropy*, 17(4):2459–2543.

Laudan, L. (1978). From theories to research traditions. In *Progress and its Problems*, chapter 3. University of California Press.

Lehrer, K. (1993). Carnap and Reichenbach on probability with Neurath the winner. In *Scientific Philosophy: Origins and Developments*, volume 1 of *Vienna Circle Institute Yearbook*. Springer.

Leitgeb, H. (2011). New life for Carnap's Aufbau? *Synthese*, 180:265–299.

Leitgeb, H. (2013). Scientific philosophy, mathematical philosophy, and all that. *Metaphilosophy*, 44(3):267–275.

Leitgeb, H. (forthcoming). Probability in logic. In Hájek, A. and Hitchcock, C., editors, *The Oxford handbook of probability and philosophy*. Oxford University Press.

Leitgeb, H. and Pettigrew, R. (2010). An objective justification of Bayesianism i: Measuring inaccuracy. *Philosophy of Science*, 77:201–235.

Lutz, S. (2012). Artificial language philosophy of science. *European Journal for Philosophy of Science*, 2:181–203.

Maher, P. (2002). Joyce's argument for probabilism. *Philosophy of Science*, 69:73–81.

Maher, P. (2007). Explication defended. *Studia Logica*, 86:331–341.

Maher, P. (2010). Explication of inductive probability. *Journal of Philosophical Logic*, 39:593–616. 10.1007/s10992-010-9144-4.

McCullagh, P. (2002). What is a statistical model? *The Annals of Statistics*, 30(5):1225–1310.

Michalos, A. (1971). *The Popper-Carnap Controversy*. Springer.

Niiniluoto, I. (2011). The development of the Hintikka program. *Handbook of the History of LogicI*, 10:311.

Olsson, E. (2014). Gettier and the method of explication: a 60 year old solution to a 50 year old problem. *Philosophical Studies*, pages 1–16.

Palmgren, E. (2009). Semantics of intuitionistic propositional logic. Lecture notes.

Parikh, R. (1995). Logical omniscience. In Leivant, D., editor, *Logic and Computational Complexity*, pages 22–29. Springer.

Paris, J. (1994). *The Uncertain Reasoner's Companion*. Cambridge University Press.

Paris, J. and Vencovská, A. (2015). *Pure Inductive Logic*. Cambridge University Press.

Peregrin, J. (2015). Carnap's inferentialism. In Schuster, R., editor, *The Vienna Circle in Czechoslovakia*. Springer.

Pettigrew, R. (2013). A new epistemic utility argument for the principal principle. *Episteme*, 10(1):19–35.

Pettigrew, R. (forthcoming). *Accuracy and the Laws of Credence*. Oxford University Press. Quotations from draft manuscript.

Polya, G. (1954). *Patterns of Plausible Inference*, volume 2 of *Mathematics and Plausible Reasoning*. Princeton University Press.

Popper, K. (1959). *The Logic of Scientific Discovery*. Routledge.

Popper, K. (1962). *Science: Conjectures and Refutaions*. New York London.

Popper, K. (1968). Theories, experience and probabilistic intuitions. In Lakatos, I., editor, *The Problem of inductive logic*, pages 285–303. North-Holland.

Potochnik, A. and Yap, A. (2006). Revisiting Galison's 'Aufbau/Bauhaus' in light of Neurath's philosophical projects. *Studies in History and Philosophy of Science Part A*, 37(3):469–488.

Predd, J. B., Seiringer, R., Lieb, E. H., Osherson, D. N., Poor, V., and Kulkarni, S. R. (2009). Probabilistic coherence and proper scoring rules. *IEEE Transactions on information theory*, 55(10):4786–4792.

Psillos, S. (2000). Carnap, the Ramsey sentence and realistic empiricism. *Erkenntnis*, 52(2):253–279.

Putnam, H. (1963). "Degree of confirmation" and inductive logic. In Schilpp, P., editor, *The Philosophy of Rudolf Carnap*, pages 761–783. Open Court publishing co.

Ramsey, F. (1931). Truth and probability (1926). *The foundations of mathematics and other logical essays*, pages 156–198.

Restall, G. (2002). Carnap's tolerance, meaning, and logical pluralism. *The Journal of Philosophy*, 99(8):426–443.

Richardson, A. (1998). *Carnap's construction of the world: The Aufbau and the emergence of logical empiricism*. Cambridge University Press.

Romeijn, J.-W. (2005). *Bayesian Inductive Logic*. PhD thesis, University of Groningen.

Romeijn, J.-W. (2011). Inductive logic and statistics. In Woods, J., Hartmann, S., and Gabbay, D. M., editors, *Inductive Logic*, volume 10 of *Handbook of the History of Logic*, pages 625–650. North-Holland.

Salmon, W. C. (1967). Carnap's inductive logic. *The Journal of Philosophy*, 64(21):725–739.

Salmon, W. C. (1991). Hans Reichenbach's vindication of induction. *Erkenntnis*, 35(1-3):99–122.

Sarkar, S. (2006). Rudolf Carnap. In Pfeifer, J. and Sarkar, S., editors, *The Philosophy of Science: An Encyclopedia*, volume 1. Routledge.

Schick, F. (1986). Dutch bookies and money pumps. *Journal of Philosophy*, 83(2):112–119.

Shimony, A. (1985). The status of the principle of maximum entropy. *Synthese*, pages 35–53.

Shimony, A. and Carnap, R. (1969). Correspondence between Carnap and Shimony. Box 1, Folder 3 Abner Shimony Papers, 1947-2009, ASP.2009.02, Archives of Scientific Philosophy, Special Collections Department, University of Pittsburgh. http://digital.library.pitt.edu/u/ulsmanuscripts/pdf/31735033440151.pdf.

Skyrms, B. (1985). Maximum entropy inferences as a special case of conditionalisation. *Synthese*, pages 55–74.

Skyrms, B. (1987). Updating, supposing and MAXENT. *Theory and Decision*, 22(3):225–246.

Skyrms, B. (1996). Carnapian inductive logic and Bayesian statistics. *Lecture Notes-Monograph Series*, 30:321–336.

Steele, K. and Werndl, C. (2013). Climate models, calibration and confirmation. *British Journal for the Philosophy of Science*, 64:609–635.

Stein, H. (1992). Was Carnap entirely wrong, after all? *Synthese*, 93(1-2):275–295.

Strawson, P. (1963). Carnap's views on constructed systems versus natural languages in analytic philosophy. In *The Philosophy of Rudolf Carnap*, pages 503–519. Open Court.

Toulmin, S. (1953). Carnap, R: Logical foundations of probability. *Mind*, 62(245):86–99.

Uebel, T. (2012). Carnap's logic of science and personal probability. In Dieks, D., Gonzalez, W. J., Hartmann, S., Staltzner, M., and Weber, M., editors, *Probabilities, Laws, and Structures*, volume 3 of *The Philosophy of Science in a European Perspective*, pages 469–479. Springer Netherlands.

Van Fraasen, B. C. (1990). Figures in a probability landscape. In Dunn, J. M. and Gupta, A., editors, *Truth or Consequences*, pages 345–356. Springer Netherlands.

Van Horn, K. (2003). Constructing a logic of plausible inference: a guide to Cox's theorem. *International Journal of Approximate Reasoning*, 34(1):3–24.

Wagner, P. (2011). Carnap's theories of confirmation. In Dieks, D., Gonzalez, W. J., Hartmann, S., Uebel, T., and Weber, M., editors, *Explanation, Prediction, and Confirmation*, volume 2 of *The Philosophy of Science in a European Perspective*, pages 477–486. Springer Netherlands.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.

Weatherson, B. (2003). From classical to intuitionistic probability. *Notre Dame Journal of Formal Logic*, 44(2):111–123.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press.

Williamson, J. (2011). Objective Bayesianism, Bayesian conditionalisation and voluntarism. *Synthese*, 178:67–85.

Williamson, J. (2013). From bayesian episemology to inductive logic. *Journal of Applied Logic*, 11:468–486.

Williamson, J. (2015). Lectures on Inductive Logic. manuscript.

Williamson, T. (2002). *Knowledge and its Limits*. Oxford University Press.

Yap, A. (2010). Feminism and Carnap's principle of tolerance. *Hypatia*, 25(2):437–454.

Zabell, S. L. (1982). W. E. Johnson's "sufficientness" postulate. *The Annals of Statistics*, 10(4):1090–1099.

Zabell, S. L. (2011). Carnap and the logic of inductive inference. In Woods, J., Hartmann, S., and Gabbay, D. M., editors, *Inductive Logic*, volume 10 of *Handbook of the History of Logic*, pages 265–309. North-Holland.