A Hybrid Theory of Induction

Adrià Segarra Torné

UNIVERSITY OF CAMBRIDGE

St John's College

This thesis is submitted for
the degree of Doctor of Philosophy

April 2023

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text. This thesis does not exceed the word limit of 80,000 words set by the Degree Committee of the Department of History and Philosophy of Science, University of Cambridge.

Abstract

A Hybrid Theory of Induction

Adrià Segarra Torné

In this thesis I motivate and develop a Hybrid Theory of Induction (HTI), and I explore some of its virtues and implications. The HTI is a hybrid second-order model of inductive support. It is a *hybrid* model of inductive support because it holds that two ingredients play a necessary role in understanding inductive support: rules and facts. It is a *second-order* model of inductive support because it is a model within which first-order models of inductive support (i.e. logics of induction) can fit. In chapter 1 I argue that we need both rules and facts to play a role in a successful account of inductive support. Rules of induction accurately describe relations of inductive support when they are warranted; facts do the warranting work. I call this type of warrant "factual warrant". The resulting account is both functional and accurate, it helps us make sense of how different rules of induction can coexist and it allows us to resolve some current debates in induction. For the purposes of chapter 1 I adopt an existing binary account of factual warrant. In chapter 2 I develop a Graded account of Factual Warrant (GFW), according to which factual warrant comes in degrees. I integrate the GFW in the HTI. I then show that the GFW illuminates the connection between factual warrant and inductive support, and it can successfully account for the role of idealisations and theory in our understanding of inductive support. In chapter 3 I argue that the HTI is also useful for agents, since it can provide methodological guidance to ensure strong inferences and conceptual guidance to assess the strength of our inferences. Finally, in chapter 4, I explore Bayesian inductive logics from the perspective of the HTI. This analysis brings to light the central role that probability models play in Bayesian inductive logics, offering a logical underpinning for some recent suggestions in Bayesian epistemology. Furthermore, throughout this thesis I analyse in detail three rules of induction from the perspective of the HTI: enumerative induction in chapter 2, causal inference in chapter 3 and Bayesian inductive logics in chapter 4. These analyses illustrate how the HTI can help us think more clearly about rules of induction, offering new tools to tackle existing challenges.

Acknowledgements

First and foremost, I want to express my deepest appreciation to my supervisor, Jacob Stegenga, for teaching me so much about what it means to do philosophy well. His unwavering support and guidance have been invaluable in shaping my research and helping me grow as a scholar.

Many others have helped me throughout this journey. I owe a special debt to André Bazzoni for his encouragement and his profound belief in my abilities. Thanks to Albert Solé and Carl Hoefer for cultivating my interest in the philosophy of science, and to Esa Díaz-León for her patience and her valuable advice. I am very grateful to John D. Norton for his helpful feedback and support. Many thanks to my advisor, Alexander Bird, and to Tim Lewens, for carefully reading excerpts of this thesis and providing valuable feedback. Thanks also to Jan-Willem Romeijn, Jonah Schupbach, Bryan Roberts, Peter Vickers, Robert Northcott, Richard Pettigrew, John P. McCaskey and Jonathan Fuller for helpful comments and suggestions.

At the University of Cambridge I have benefited from being immersed in a stimulating research environment. I owe an enormous debt of gratitude to Cristian Larroulet Philippi, for his invaluable contributions to my work and for offering me his friendship. I thank everyone in the Department of History and Philosophy of Science for many interesting exchanges during seminars, reading groups and pub nights. A very special thanks to Ahmad Elabbar and Miguel Ohnesorge, good friends from whom I have learned so much. I have also been part of an outstanding group of graduate students. I'm grateful to Benjamin Chin-Yee, particularly for introducing me to the Cambridge wine scene. Thanks also to Adrian Erasmus, Zinhle Mncube, Hamed Tabatabaei Ghomi, Oliver Holdsworth, Sophia Crüwell, Ina Jäntgen, Charlotte Zemmel, Matthew Gummess, Dominique Waissbluth-Kingma, Michaela Egli, Arthur Harris and Claudio Davini for their helpful feedback and camaraderie. At Cambridge, I had the great pleasure of being a member of St John's College. I thank everyone at the Lady Margaret Boat Club and the Redboys for keeping me sane and healthy over the past few years. I'm grateful to all the staff at John's, specially to Farida, Norma, Anna (both of them), Cat, Gabi, Aida, Olivia, John, Nick, Seb and Storm, for their immense kindness. I thank my College tutors, Morag Morrison-Helme and Amanda Sferruzzi-Perri, for their relentless help and guidance. I would also like to acknowledge the generous financial support provided by the Cambridge Trust, the Department of History and Philosophy of Science, and St John's College, without which this research would not have been possible.

On a more personal note, I want to express my deepest gratitude to my parents, Joana and Xavier, and to my brother, Xavi. I wouldn't be here today if it weren't for their unconditional love and support. I'm grateful to my friends in Catalonia, for all the joy they bring to my life. Finally, to Laura, thank you for making my days brighter with your love and reassurance.

Contents

Introduction

Mature science is an extremely successful enterprise. Scientists rely on observed evidence to establish powerful and fruitful theories, which help us understand our world and predict incredibly well how it will behave in the future. It is a distinctive feature of scientific hypotheses that they can receive support from empirical evidence. Eddingtons's measurement of the gravitational deflection of starlight passing near the sun supported Einstein's general theory of relativity. Darwin's observation of patterns in the distribution and features of organisms supported his theory of evolution. However, these relations of support between evidence and hypotheses remain elusive. In this thesis I explore how we should think about them.

Relations of support between evidence and hypotheses are quite special. In order to learn about the world, scientific hypotheses must go beyond the evidence in one way or another. We observe a few samples of radium chloride and they all have the same crystallographic form; we might then hypothesize that *all* radium chloride forms crystals just like the ones we have observed. Our observation supports this hypothesis, but it does not entail it. The fact that all observed samples of radium chloride have the same crystallographic form is compatible with the next sample of radium chloride forming a different crystal. This relation of support, which is weaker than entailment, is known as *inductive support*, and the study of such relations is the business of *inductive logic*.

The philosophical study of induction goes back, at least, to Aristotle. Cicero coined the Latin term *inductio* to translate Aristotle's *epagoge*, a concept that Socrates allegedly used in his discussions (McCaskey, 2020, 2022). Since then, there have been two important traditions in the study of induction. The rule-based tradition, mirrored on the study of deduction, assumes that induction is governed by some underlying rules. According to this tradition, once the right rules are exposed they should be able to inform us about inductive support and guide us in our inductive practices. The material tradition, instead, assumes that the degree of support that a piece of evidence yields a hypothesis does not depend on any underlying rules, but on matters of fact. In this thesis I motivate and develop a theory of induction where both rules and matters of fact play a role.

Work on induction is often tackled either from a logical or an epistemic perspective, and it will be good to disentangle both projects at the outset for clarity. Richard Royall begins his 1997 book on the concept of evidence by distinguishing the following three questions:

1. What does the present evidence say?

2. What should you believe?

3. What should you do?

Question (1) is the business of inductive logic, while question (2) is the business of epistemology. Crucially, as Sober points out (Sober, 2008, pp. 3–4), answering question (2) requires more than an answer to (1) (and answering question (3) requires more than an answer to (2), but that is irrelevant for our current purposes). Sober provides the following example to illustrate the difference between answering question (1) and answering question (2):

> Suppose you are a physician and you are talking to the patient in your office about the result of his tuberculosis test. The report from the lab says "positive." This is your present evidence. Should you conclude that the patient has tuberculosis? You want to take the lab report into account, but you have other information besides. For example, you previously had conducted a physical exam. Before you looked at the test report, you had some opinion about whether your patient has tuberculosis. The lab report may modify how certain you are about this. You update your degree of belief by integrating the new evidence with your prior information. (Sober, 2008, p. 4)

In this case, the lab report may support the hypothesis that the patient has tuberculosis, but that does not determine what you should believe about whether the patient has tuberculosis. Further considerations, like your prior beliefs and other sources of evidence, may influence what you should believe in this scenario.

Gilbert Harman's influential work may also help us grasp the distinction between logic and epistemology more clearly. Harman contrasts deductive logic with what he calls "theories of reasoning". He argues that both endeavours are often mistakenly conflated, because the term "logic" has been used inconsistently in the literature. Harman tells us that the subject matter of theories of reasoning are the psychological events or processes that constitute reasoning, hence in the realm of epistemology, while the subject matter of deductive logic are relations between propositions (Harman, 2009). Thus, he reaches the following conclusion:

> Logical principles are not directly rules of *belief revision*. They are not particularly about belief at all. For example, *modus ponens* does not say that, if one believes *p* and also believes *if p then q*, one may also believe *q*. Nor are there any principles of belief revision that directly correspond to logical principles like *modus ponens*. Logical principles hold universally, without exception, whereas the corresponding principles of belief revision would be at best *prima facie* principles, which do not always hold. It is not always true that, if one believes *p* and believes *if p then q*, one may infer *q*. The proposition *q* may be absurd or otherwise unacceptable in the light of one's beliefs, so that one should give up either one's belief in *p* or one's believe in *if p then q* rather than believe *q*. And, even if *q* is not absurd and is not in conflict with one's other beliefs, there may simply be no point adding it to one's beliefs. The mind is finite. One does not want to clutter it with trivialities. It would be irrational to fill one's memory with as many as possible of the logical consequences of one's beliefs. That would be a terrible waste of time, leaving no room for other things. (Harman, 1984, p. 107–108, emphasis in original)

Both Sober and Harman's observations point towards the idea that, while logical relations between proposition influence what we should believe, they do not fully determine it. Further considerations, beyond the logical relations of support between propositions, must be considered in determining what we should believe, as well as what we actually believe. Thus, there is a gap between logic and reasoning, and the challenge of bridging this gap is known as "Harman's challenge" (Steinberger, 2022). This is a particularly clear way to see that logic and epistemology are distinct endeavours.

Elliot Sober offers yet another perspective to distinguish logic and epistemology. In the following fragment, Sober contrasts the concept of "evidence" with the concepts of "acceptance" and "rejection". The concept of "evidence" plays a central role in inductive logic, while the concepts of "acceptance" or "rejection" play a central role in epistemology. In contrasting these concepts, Sober rightly notices the following:

> It is interesting that the concept of evidence relates pairs of propositions to each other, while the concepts of acceptance and rejection relate propositions to persons. Smoke is evidence for fire, regardless of whether any agent takes this fact to heart. However, rational acceptance (or rejection) means that a person is justified in accepting (or rejecting) some proposition. The present disciplinary divide between philosophers of science and epistemologists coincides to a considerable degree with this distinction between questions concerning how propositions are related to each other and questions concerning how propositions are related to persons. (Sober, 2008, p. 6)

I believe this distinction is exactly right. Inductive logic is concerned with how propositions relate to other propositions, independently of whether these propositions are believed by some agent. Epistemology is concerned with how agents relate to propositions (by accepting, rejecting, or believing them to a certain degree, for example). Relations between propositions may influence how agents relate to such propositions, but they do not fully determine it, as Sober and Harman point out in the fragments above.

In conclusion: logic, on the one hand, explores the relations of support that arise between *propositions* – for example, between the proposition expressing the evidence available and the proposition expressing the hypothesis of interest. Epistemology, on the other hand, explores how we relate to these propositions (the doxastic attitudes we have towards them). Logic and epistemology are closely related, but distinct (de Grefte, 2020; Harman, 1984, 2002). Therefore, we should avoid confusing or conflating these disciplines. This dissertation is primarily concerned with logic, although chapters 3 and 4 also explore some epistemic consequences of the perspectives on logic introduced in the previous chapters.

Before summarizing the work I will do in this thesis, let me point to some of the work that I will *not* do. First, I will not discuss Hume's Problem of Induction. Hume's Problem is the problem of justifying our inductive inferences. It is an important challenge, and it must be addressed. A lot has been written about it, and I point the reader to (Schurz, 2019) or (Norton, 2021b) for some recent literature on this problem. However, Hume's Problem is an epistemological problem (de Grefte, 2020, §4), and my concern in this thesis is with inductive logic. Since I do not pretend to argue that my views on inductive logic can resolve Hume's Problem, I will not talk about it. Second, I will not commit to any metaphysics of inductive support. I tend to think of inductive support in terms of possible worlds, similarly to Carnap. According to Carnap, to say that a proposition E confirms a proposition H (to a certain degree) is to say something about the relation between $\mathfrak{R}(E)$ and $\mathfrak{R}(H)$, where $\mathfrak{R}(\cdot)$ is the range of a proposition, that is, the set of possible worlds where this proposition is true. This is still an influential view today (e.g. Hawthorne, 2018). However, in this thesis I present my views in metaphysically neutral terms. We can explore the problems with existing accounts of inductive support, and suggest ways to address these problems, without committing to any metaphysics of this concept.

I am now ready to summarize the work that I will do. In this thesis I will develop a Hybrid Theory of Induction (HTI) and explore some of its virtues and implications. A central motivation for the HTI is the observation that rules of induction, which describe relations of inductive support, sometimes fail at this task. The rule of enumerative induction, for example, might tell us that "an instance confirms the generalization". This particular rule tells us that inductive support arises between two propositions when (roughly speaking) one is an instance of the other – hence, it is providing a description of a specific relation of inductive support. However, this description is not accurate, since it is not always the case that an instance confirms the generalization. While the proposition that "this stone falls when dropped" seems to support the proposition that "all stones fall when dropped", it is not the case that the proposition "Bill wears a grey shirt today" supports the proposition that "Bill wears a grey shirt every day". Thus, the rule of enumerative induction is not universal, since it does not always provide an accurate description of the relations of inductive support that arise between instances and generalizations. This turns out to be a problem for any rule of induction: none of them is universal. However, this problem can be solved by realising that all rules of induction seem to work *somewhere*, and we can delimit the domain where a given rule of induction provides adequate descriptions of inductive support. The rule of enumerative induction introduced above seems to work when the properties being generalized are uniform: on the one hand, stones are always attracted by the gravity of the earth, but, on the other hand, Bill changes his clothes often. In short, the HTI tells us that rules of induction provide adequate descriptions of inductive support when they are applied in the right domains, domains in which the right facts about the matter of the induction obtain. Thus, according to the HTI, we need to appeal to both rules and matters of fact in order to properly describe relations of inductive support. This need to appeal to two ingredients in order to understand relations of inductive support is what makes the HTI *hybrid*. Next, I provide a brief summary for each chapter of the thesis.

Chapter 1. The Need for Rules and Facts

In chapter 1 I argue that a successful model of inductive support requires two ingredients: rules and facts.

According to the rule-based tradition in the philosophy of induction, which has dominated for the last couple of centuries, inductive arguments are warranted by rules. Exposing the right rules of induction has been an important task for philosophers and logicians alike, and there is no shortage of candidates nowadays, like Bayesianism or Inference to the Best Explanation. Rules of induction provide functional accounts of inductive support (they provide information about inductive support), but no rule is universal; hence, no rule is by itself an accurate model of inductive support.

According to the material tradition in the philosophy of induction, inductive arguments are not warranted by rules but by matters of fact. Norton's Material Theory of Induction (MTI) is an influential view within this tradition. Norton's MTI fails to provide an account of inductive support (it does not provide information about inductive support), but it provides a good account of inductive warrant that can help us define the domain of validity of each rule of induction. The warrant of an inductive argument is the source of inductive support, and Norton's MTI tells us that this source is constituted by facts about the matter of the induction.

Despite their limitations, both approaches to induction are illuminating important aspects of how relations of inductive support work. In this chapter I introduce the foundations of the HTI, by motivating the need for both rules and matters of fact in our understanding of inductive support. According to the HTI, rules of induction accurately describe relations of inductive support when they are warranted, and facts do the warranting work. Crucially, the HTI allows us to address the main challenges that each tradition faces while retaining their strengths, thus obtaining a functional and accurate account of inductive support. The HTI also provides a useful general framework to examine and understand induction, to make sense of how different rules of induction can coexist and to tackle some problems in epistemology. Moreover, the HTI allows us to clarify and resolve some current debates on induction, like the apparent divide between Norton and rule-based theorists.

Chapter 2. Logical Virtues of a Graded Account of Factual Warrant

In chapter 2 I refine the HTI, which was introduced in chapter 1, by supplementing it with a Graded Account of Factual Warrant (GFW). I then highlight two logical virtues of the GFW. In this chapter I provide an analysis of the rule of enumerative induction in terms of the HTI and I rely on this analysis to make my points.

Rule-based accounts of induction have dominated the literature for many years, so their underlying rules have been explored and discussed in detail. In this chapter I briefly introduce my understanding of rules of induction as sentence schemas. However, I focus on the more neglected task of developing the role of facts in understanding inductive support. Although the idea that matters of fact may play a role in understanding inductive support is not new, Norton's MTI is the only contemporary well-developed account of factual warrant available. While, for the purposes of chapter 1, I embraced Norton's MTI, in this chapter I develop a different Graded account of Factual Warrant (GFW). The crucial difference between Norton's MTI and the GFW is that, while Norton's MTI is a binary account of warrant (according to which arguments are either warranted or they are not), according to the GFW factual warrant comes in degrees. In this chapter I introduce this account of warrant and integrate it into the HTI, thus obtaining a fully developed version of the HTI.

This chapter has three goals. First, and foremost, to introduce the GFW and the resulting version of the HTI. Second, to provide two arguments in favour of the GFW and the updated HTI. This will be accomplished by highlighting two logical virtues of the GFW, in which the MTI fails: (1) the GFW illuminates the connection between factual warrant and inductive support, and (2) the GFW can successfully account for the role of idealisations and theory in our understanding of inductive support. The third and last goal is to illustrate the GFW and the resulting version of the HTI. This will be accomplished by analysing the rule of enumerative induction in terms of the HTI. This rule and the resulting analysis will be used as a running example throughout this chapter. Enumerative induction is a simple and popular rule, and, despite its limitations, it is widely used both in everyday and scientific reasoning. Therefore, it will be interesting to tackle this rule of induction first. In chapters 3 and 4 I explore other rules of induction.

Chapter 3. Epistemic Virtues of a Graded Account of Factual Warrant

In chapter 3 I provide an analysis of a rule used for causal inferences in the context of comparative group studies, like RCTs, the merits of which have been widely debated. I then use this example to highlight two epistemic virtues of the GFW. While in the previous chapters I focus on inductive *logic* and, accordingly, I talk about *arguments*, in this chapter I extract some epistemic consequences of the HTI and, accordingly, I talk about *inferences*.

I start by exposing a particular rule of induction used for causal inference in the context of comparative group studies. This analysis shows that the HTI brings clarity to debates around rules of induction. The HTI forces us to express our rule of interest precisely, in the form of a sentence schema, and to expose the ideal conditions under which this rule provides perfectly accurate information on inductive support. I call these "ideal warranting conditions". By analysing rules of induction in these terms, we can disentangle similar but different rules that would otherwise fall under the same slogan, and clarify the conditions under which each rule is appropriate.

I then show why this analysis is useful for the agents wishing to make inferences following any given rule of induction. The ideal warranting conditions of a rule can function as a methodological guide for researchers to ensure strong inferences. I develop this point and provide an argument in favour of random subject allocation in comparative group studies. Against some popular critiques, I show that randomization does indeed take us closer to the ideal warranting conditions for a certain type of causal inference. Furthermore, the ideal warranting conditions of a rule also function as a conceptual guide to those features of the world that are relevant in assessing the strength of an inductive inference. Thus, analysing rules of induction in terms of the HTI is useful for agents in at least two ways: the analysis provides a methodological guide which can help us ensure strong inferences, and it also provides a conceptual guide to assess the strength of our inferences.

Chapter 4. Bayesian Inductive Logics and the HTI

In chapter 4 I explore Bayesian inductive logics from the perspective of the HTI. In doing so, I hope to illustrate three points.

First, that the HTI provides a unified framework to situate many open debates and disagreements about Bayesian inductive logics. In order to articulate a Bayesian inductive logic in the form of a sentence schema, as the HTI requires, one must take a stance with respect to many open debates. We must be explicit about our understanding of confirmation, whether (and how) we choose to quantify it, or what axioms of probability we prefer, for example. Without these commitments our position is not fully specified, and it cannot be articulated in the form of a sentence schema.

The HTI provides a clear framework to understand which questions need to be answered, and how our answers impact the resulting logic.

Second, I will make explicit the central role that probability models play in any Bayesian inductive logic. In particular, I will bring to light the relation between the accuracy of our probability models and the accuracy of the information on inductive support that we obtain by using them in any Bayesian inductive logic. Any Bayesian rule must articulate the notion of confirmation in terms of probabilities. These probabilities are given by, or calculated from, a probability model of our system of interest. Hence, the ideal warranting conditions for any Bayesian rule will require that the probability model being used is a perfectly accurate model of our target system. Under these ideal conditions, the information on inductive support that our Bayesian rule provides is true; as our probability model becomes less accurate, the information on confirmation provided by our Bayesian rule becomes less truthlike.

Last, I will show that understanding Bayesian inductive logics from the perspective of the HTI has some epistemic consequences too. The ideal warranting conditions for Bayesian inductive logics function as a methodological guide in the following sense: they tell us that we should maximize the accuracy of our models, since more accurate models will yield more truthlike information on inductive support, which will better inform our inferences. These ideal warranting conditions also function as a conceptual guide to the factors that are relevant in assessing the strength of our inferences: these will be the factors which are relevant in assessing the truthlikeness of the ideal warranting conditions for our Bayesian rule, hence, the factors which are relevant for the accuracy of our probability models.

These epistemic consequences are in line with some recent suggestions in the literature on Bayesian epistemology. Morey, Romeijn, and Rouder (2013) and Gelman and Shalizi (2013), for example, urge that philosophical accounts of Bayesianism should account for the practice of model-checking (that is, checking one's own models). The need for model-checking follows from looking at Bayesianism from the perspective of the HTI: we must check our models in order to assess their accuracy, so that we can assess the strength of our inferences. Relatedly, Sprenger and Hartmann (2019, pp. 311–26) suggest a particular understanding of Bayesianism as a form of model-based reasoning, consistent with the epistemic consequences discussed in this chapter. Thus, the HTI offers a logical underpinning for these epistemic positions. Furthermore, by making the ideal warranting conditions of Bayesian rules explicit, the HTI can help us develop better bridge principles to understand the connection between Bayesian inductive logics and Bayesian epistemologies. In this chapter I illustrate how we can go about this, by showing how the HTI allows us to generalize a specific bridge principle for Bayesian inductive logics suggested by Fitelson (2006).

Chapter 1

The Need for Rules and Facts

1.1 Introduction

What determines whether the premises of an inductive argument support its conclusion? Traditionally, there have been two types of answers to this question. Some philosophers believe that *rules* of induction tell us whether the premises P support the conclusion C in an inductive argument (and, maybe, to which degree). Huge efforts have been directed at exposing the right rules of induction. This view is still prevalent nowadays, with Bayesianism being the most popular theory of induction in this category (e.g. Hawthorne, 2018). Other philosophers, however, have argued that it is not rules but *matters of fact* that determine whether P supports C in an inductive argument. This approach, although less popular today, has been present since antiquity (McCaskey, 2020). Norton's Material Theory of Induction (MTI) (Norton, 2003, 2021c) is a prominent contemporary theory of induction in this category. Both of these traditions face several challenges, but both are getting at something important about how induction works. In this chapter I will introduce a Hybrid Theory of Induction (HTI), an account which acknowledges the role of both *rules* and *matters of fact* in a successful theory of induction. Hopefully, the HTI will articulate a view that is actually shared by most agents in this debate, thus providing a common framework to understand and tackle disagreements about induction with more clarity.

The aim of this chapter, however, is not to provide a fully developed version of the HTI. In this chapter I show that we need both rules and matters of fact to play a role in our understanding of inductive support, and I sketch what these roles are. In order to achieve this aim, I will temporarily adopt Norton's account of the role of matters of fact in understanding inductive support (§1.4.2). This will suffice to illustrate the basic idea behind the HTI: that we need rules to play a role if we want our account of inductive support to be functional, and we need facts to play a role if we want our account of inductive support to be accurate. Later on, in chapter 2, I will develop a more refined view of the role of facts in understanding inductive support, and I will integrate it in the HTI. Thus, the HTI will not be fully developed until next chapter.

Some preliminary clarifications are due. I take an "argument" to be a system of two propositions, a premise P and a conclusion C (both of which may be conjunctions of more simple propositions), where P is supposed to provide support for C. We should distinguish between "arguments", which are composed of propositions, and "inferences", which are composed of beliefs. Keeping this distinction clear is important to avoid conflating logic and epistemology. While these terminological choices are popular (e.g. Harman, 2002; Hawthorne, 2018), the usage of the terms "argument" and "inference" is inconsistent in the literature. Notably, Norton chooses to use the term "inference" to refer to what I have called "argument" above (e.g. Norton, 2021c, p. 20), a choice that may cause confusion in some readers. I will use the terms "induction" or "inductive argument" to refer to ampliative arguments, that is, arguments in which the conclusion contains information not contained in the premises.[1] In an inductive argument, P supports C but does not entail it. In such an argument, the relation of support between P and C is often referred to as "inductive support" (or "confirmation", a popular term in classical texts), and the strength of that relation as "inductive strength", "inductive degree of support" or simply "degree of support".[2] I will use the term "warrant" to refer to the ultimate reason why P supports C in an argument – in other words, the warrant of an argument is the source of support.[3] The warrant of a deductive argument is its form, because that is the reason why P supports C in that argument; it is the source of support. In the same way, the warrant of an inductive argument is the source of inductive support, but the nature of inductive warrant is a contentious matter (see §1.3). The HTI is a model of inductive support. I take the two main desiderata for such a model to be "functionality" and "accuracy". A functional model of inductive support must actually inform us about the relations of support in inductive arguments; an accurate model must do this correctly.

A last and crucially important clarification: the HTI is concerned with inductive *logic*, not with epistemology. Logic and epistemology are sometimes conflated or confused. This confusion is caused, in part, because the term "logic" is used in several ways in the literature (Harman, 1984, 2002). Sometimes "logic" is used to mean a "theory of reasoning" or a "theory of rational belief". This is not how I will use the term. By "logic" I mean a "theory of the relations of support between propositions", not between the beliefs of an agent. *Inductive* logic, in particular, studies the relations of support between propositions in ampliative arguments (Hawthorne, 2018; Norton, 2021c, p. 52). Let me illustrate this distinction by focusing on argument [Salt].

$$\begin{array}{ll} P: & \text{This sample of radium chloride is monoclinic} \\ C: & \text{All samples of radium chloride are monoclinic} \end{array} \qquad \text{[Salt]}$$

According to Norton (2021c, p. 43), P supports C in argument [Salt] because the fact that "generally, each crystalline substance has a single characteristic crystallographic form" obtains. Thus, this fact is the warrant of argument [Salt]. He calls this fact the "Weakened Haüy's Principle". Norton's view is about inductive logic, since it concerns the relation of support between both

[1] Notice that this choice makes abduction a type of induction. Some philosophers define induction more narrowly. However, this choice is irrelevant for the view I present here – the HTI is a general thesis about all ampliative arguments, whether we like to call them inductive or not.

[2] Sometimes the term "evidential support" is used in a logical sense instead of "inductive support" (e.g. Hawthorne, 2018), but I will avoid this usage, since some philosophers use the term "evidential support" in an epistemic sense (e.g. Fitelson, 2006).

[3] In this terminological choice I am following Norton. While he is not explicit in his usage of the term warrant, the understanding of this term as referring to the source of inductive support becomes evident in several of his writings (e.g. 2021c, pp. 61–62; 2021a, pp. 122–23). Other authors use the term "warrant" in a wider sense (e.g. Reiss, 2020, p. 14), and this can cause confusion. The confusion is aggravated by the fact that most authors leave the term "warrant" undefined in their writings, and that the term is sometimes used in a logical sense and sometimes in an epistemic sense.

propositions in an ampliative argument. It is a different matter, however, to ask when an agent is justified in making this inference. This is an epistemic question. An externalist about epistemic justification who agreed with Norton's views on inductive logic would probably tell us that an agent is always justified in making this inference, as long as the Weakened Haüy's Principle obtains; the internalist counterpart, however, would probably require the agent to know (or be justified in believing) the Weakened Haüy's Principle. Both epistemic views are compatible with the same logical picture. Too often the boundaries between logic and epistemology are blurred, and this leads to misunderstandings.

This is not to say, of course, that logic and epistemology are independent endeavours, since they are closely related in several ways (de Grefte, 2020). The relations of support between our beliefs depend, in part, on the relations of support between the propositional content of those beliefs, something that our inductive logic is supposed to capture. Exploring the epistemic consequences of the HTI will surely come with its own challenges. I explore some of these epistemic consequences in chapters 3 and 4.

Here is how I will proceed. In section 1.2 I introduce an important distinction between inductive *support* and inductive *warrant*. In section 1.3 I present different ways in which theories of induction can be classified, and I choose to classify them according to whether or not they rely on rules to warrant inductive arguments. In section 1.4 I introduce several theories of induction following this classification. In section 1.5 I argue that "rules of induction" must play a role in our model of inductive support if it is to be functional. In section 1.6 I argue that "local facts" must play a role in our model of inductive support if it is to be accurate. In section 1.7 I introduce the HTI, which embraces the lessons from sections 1.5 and 1.6. In section 1.8 I argue that the HTI is not a completely new idea, but an articulation of a view that has long been present in the literature but has not been made explicit. Section 1.9 summarises some advantages of the HTI. Section 1.10 concludes.

1.2 Inductive Support and Inductive Warrant

Throughout this dissertation I will use the concepts of "inductive support" and "inductive warrant" extensively. Thus, it will be good to clarify my understanding of these notions. This is the purpose of the current section.

Relations of inductive support are logical relations of support between propositions in ampliative arguments. As I anticipated in the introduction to this dissertation, I will not explore the metaphysics of such relations, but it may be helpful to recall how Carnap liked to think of them. According to Carnap, to say that a proposition E confirms (or supports) a proposition H (to a certain degree) is to say something about the relation between $\mathfrak{R}(E)$ and $\mathfrak{R}(H)$, where $\mathfrak{R}(\cdot)$ is the range of a proposition, that is, the set of possible worlds where this proposition is true. This is still the standard understanding of inductive support today (Hawthorne, 2018).

I use the term "inductive warrant", following Norton, to refer to the *source* of inductive support, that is, the ultimate reason why the premise of an argument supports its conclusion to a particular degree (Norton, 2021c, pp. 61–62; 2021a, pp. 122–23). Relations of inductive support do not always arise, and when they do, they are not always equally strong. Thus, there must be something about each argument and its context that determines whether a premise supports a conclusion (to a given degree). That ultimate reason, whichever it is, is the source of inductive support, and Norton calls that source the "warrant" of the argument. Section 1.3 introduces several accounts of induction that posit different sources of inductive support: the syntax of the argument, some non-syntactic condition, the meaning of the terms in the argument, or local facts about the world, to name a few.

In order to see clearly that considerations about inductive support and considerations about inductive warrant come apart, let us focus on a specific example. Take the following argument:

P : A six-sided die has landed on an even number
C : The die has landed on 2

Let us be Bayesians for a minute and assume that we take the posterior probability of the conclusion of this argument as an adequate measure of inductive support. Let me explore two scenarios. In scenario 1, our die is fair and we model it as such. If we model this six-sided die as a fair die, the degree of support in this argument is 1/3. This is quite intuitive: given that the die has landed on an even number, and that there are three even numbers in a six-sided die, there is a probability of 1/3 that it has landed on 2.[4]

Now, one may ask what the *source* of inductive support is here; why is it the case that the degree of support in this argument is 1/3? In fact, the degree of support in this case is 1/3 if and only if our system behaves in a particular way; that is, if each side of the die has the same probability of landing on top. In other words, the degree of support in this argument is 1/3 if and only if our die is indeed a fair die. Hence, the inductive warrant in this case is the condition that our die is fair. Since this is the case in scenario 1, then the information on inductive support we have obtained is fully warranted, despite the degree of support itself being quite weak – thus, in scenario 1 we have full warrant without much inductive support.

Let's move to another scenario (scenario 2) in which our die is heavily imbalanced but we still model it as a fair die. The die in scenario 2 is such that it lands on number 2 almost every time. In fact, for this die, the probability that it lands on 2 given that it has landed on an even number is 0.99, so the degree of support in our argument of interest is very high (0.99).[5] However, we are still modelling this die as a fair die, so we obtain very inaccurate information on inductive support – information telling us that the degree of support in our argument of interest is 1/3. The warrant for this information is very low, since we are relying on a very inaccurate model of our die. Thus, in scenario 2 we have very low warrant, but very strong inductive support.

Let me restate the divide between inductive support and inductive warrant in slightly different terms. Considerations about inductive support have to do with the strength of the relation of support between the propositions in an ampliative argument; considerations about inductive warrant have to do with the accuracy of that information. For instance, our assessment of inductive

[4]We can choose any other Bayesian measure of confirmation we wish, and it will still be the case the the degree of support in this argument is short from maximal.

[5]We can choose any other Bayesian measure of confirmation we wish, and it will still be the case that the degree of support in this case is very close to maximal.

support in the previous argument is only accurate if we are using the right model for our system, that is, if that six-sided die is indeed fair. If our die is not fair, then the assessment of inductive support that we have obtained by modelling our die as a fair die will be inaccurate. We can choose to model our die in many different ways, which will be better or worse at representing our system. Each model will result in a given assessment of inductive support (these are considerations about inductive support), and each of these assessments will be more or less accurate depending on how well the corresponding model is representing our system (these are considerations about inductive warrant).

Sometimes, considerations about inductive support and considerations about inductive warrant might be difficult to tell apart. As will become evident through chapters 2 and 3, this will be the case when we focus on rules of induction that posit a relation of entailment (i.e., of maximum degree of support). However, this must not always be the case. Chapter 4 examines Bayesian inductive logics in detail, where considerations of inductive support can be more easily distinguished from considerations of inductive warrant. Those readers that are eager to see how these considerations come apart in the case of Bayesian inductive logics can jump to chapter 4 right after chapter 2.

1.3 Classifying Theories of Induction

Theories of induction are usually classified according to which type of warrant they rely on. In this section I introduce some systems of classification to prevent misunderstandings.

In classical texts it is typical to distinguish between syntactic and semantic theories of induction. Syntactic theories of induction rely on a syntactic warrant, i.e., they hold that the source of inductive support of an argument is its form, just like in deductive arguments. Hempel's model of confirmation by instances (Hempel, 1943, 1945a, 1945b) is an example of a syntactic theory of induction, and so is Carnap's syntactic approach to Bayesian logicism (Carnap, 1950/1962). According to Hempel, for example, P confirms C in argument [Swan] because this argument has the form expressed in argument [Swan Formal] and this is (according to Hempel) a valid form (Hempel, 1945a, 1945b).

$$\begin{array}{ll} P: & \text{There is a white swan} \\ C: & \text{All swans are white} \end{array} \qquad \text{[Swan]}$$

$$\begin{array}{ll} P: & \exists x(Sw(x) \wedge Wh(x)) \\ C: & \forall x(Sw(x) \rightarrow Wh(x)) \end{array} \qquad \text{[Swan Formal]}$$

In contrast, semantic theories of induction rely on a semantic warrant. These theories defend that the source of inductive support of an argument is the meaning of the terms in its sentences, not its form. Sellars' material theory of induction (1953), as well as Brandom's material theory (2000) based on Sellars', are examples of semantic theories of induction. According to Brandom, for example, P supports C in argument [Lightning] because of "the contents of the concepts *lightning* and *thunder*, as well as the temporal concepts" (Brandom, 2000, p. 52).

P :	Lightning is seen now	[Lightning]
C :	Thunder will be heard soon	

However, the syntactic/semantic system of classification is not exhaustive. Most contemporary views on Inference to the Best Explanation or Bayesianism, for instance, cannot be properly classified as syntactic or semantic.

Instead, I will distinguish between rule-based and non-rule-based theories of induction.[6] According to rule-based theories, the source of inductive support is a specific rule. I will call this type of warrant "rule-based warrant". Inference to the Best Explanation (IBE), for example, is a rule-based theory of induction in this sense. According to IBE, evidence E supports hypothesis H in an argument if (and because) H entails E and also best explains it. For instance, according to IBE, P supports C in argument [Snowshoes] because C entails P and also best explains it (Lipton, 2004, p. 56).

P :	There is a track of snowshoes in the snow	[Snowshoes]
C :	A person with snowshoes walked in the snow	

In this case, the alternative explanations for P (e.g. that a trained monkey on snowshoes walked in the snow) are worse than C, according to some account of what is an explanation and what makes it better or worse.

In contrast, according to non-rule-based theories of induction, the warrant of an inductive argument is not to be found in any rules. The classic non-rule-based approach to induction is the "material" approach, according to which the warranting work is done by the "matter of the induction" in one way or another. Although rule-based theories have dominated for the last couple of centuries, the material approach is not a new idea: "material theories of induction prevailed all through antiquity and from the Renaissance to the mid-1800s" (McCaskey, 2020, p. 1).

The central tenet of the material approach, however, has been fleshed out in different ways. Brandom's and Sellars' accounts introduced above, for example, are material theories of induction that rely on a semantic warrant. According to these accounts, the warranting work is specifically done by the meaning of the terms in the argument. A different contemporary material theory of induction is Norton's MTI (2003, 2021c). According to the MTI, the source of inductive support are local facts pertaining to the matter of the induction, so it's not "meanings" but "local facts" that are doing the warranting work here.

In order to distinguish Norton's MTI from other material theories, I will call the MTI a "factual" theory of induction, and I will call the type of warrant it relies on "factual warrant". Norton argues, for instance, that P supports C in argument [Bismuth] because the fact that "generally, chemical elements are uniform in their physical properties" obtains (Norton, 2003, p. 649).

[6] I avoid here the alternative terminologies "formal vs. non-formal" or "formal vs. material" to prevent confusions. These terms mean different things for different authors, and they are often used inconsistently. The term "formal", for example, is often used to refer to syntactic theories of induction, but Norton uses this term to refer to what I am calling "rule-based theories of induction".

P :	Some samples of the element bismuth melt at 271ºC	[Bismuth]
C :	All samples of the element bismuth melt at 271ºC	

Crucially, as we will see, classifying theories of induction as rule-based and non-rule-based will prove to be useful, since theories within the same group will share important advantages and challenges. The HTI will be a "hybrid" theory in the sense that, according to it, both rules and local facts play a necessary role in understanding inductive support.

1.4 Introducing Theories of Induction

1.4.1 Rule-based theories of induction

The HTI will maintain the idea that inductive arguments are, in first instance, governed by rules. The HTI will be compatible with different rules of induction being valid in different domains. Thus, it will suffice here to summarise the standard classification of rule-based accounts (Earman & Salmon, 1992; Norton, 2005), to get acquainted with the different kinds of rules available to us.[7]

Rule-based theories of induction are often classified in three broad families. Within each family fall several theories of induction, which rely on some development of the same underlying rule or principle. First is the family of "inductive generalization". Theories belonging to this family rely on the principle that an instance confirms the generalization. Hempel's account of confirmation by instances (Hempel, 1943) would belong in this family. Second is the family of "hypothetical induction". Theories belonging to this family rely on the principle that a hypothesis's ability to entail the evidence is a mark of truth. Accounts of IBE (e.g. Lipton, 2004) are developments of this principle and thus would belong in this family. Last, is the family of "probabilistic induction". Theories belonging to this family rely on the principle that relations of inductive support are probabilistic. Any form of Bayesianism (e.g. Hawthorne, 2018) or likelihoodism (e.g. Sober, 2008) for instance, would belong in this family.

1.4.2 Non-rule-based theories of induction

In this section I introduce Norton's MTI and the mechanism of factual warrant. While the MTI is not the only non-rule-based theory of induction available, it is a very powerful and flexible one. As already stated, the main classic non-rule-based approach to induction is the material approach. Different material theories of induction share, of course, a commitment to the idea that the warranting work is done by "matters of fact", and they articulate this commitment in different ways. According to Brandom, for instance, in argument [Lightning] the content of the concepts "lightning", "thunder" and the temporal concepts warrant the argument. This same idea can also be captured by Norton's MTI, which would tell us that we need a warranting fact of the sort "generally, after lightning is seen, thunder is heard". Despite these similarities, factual warrant still seems

[7] This standard classification is not exhaustive. It is a useful way to group and introduce some well-known rule-based theories of induction, but some rule-based approaches to induction will not fit in this system.

more flexible than semantic warrant, since factual warrant can contain information beyond that which is expressed by the terms that appear in the argument, in contrast with semantic warrant.

Norton's MTI (2021c) states that all inductive arguments are warranted by facts pertinent to the matter of the induction, to which Norton refers as "material facts" or "material postulates" – I will sometimes just call them "local facts". Hence, in an inductive argument, *P* supports *C* because a particular material fact obtains, not because the argument follows any rules or patterns. These warranting facts are domain-specific, so there are no universal warrants. Norton's view is summarised in his lemma, "all induction is local".

Let's illustrate the MTI by comparing arguments [Bismuth] and [Wax] (Norton, 2003, p. 649).

$$\begin{array}{rl} P: & \text{Some samples of wax melt at 91ºC} \\ C: & \text{All samples of wax melt at 91ºC} \end{array} \qquad \text{[Wax]}$$

These arguments have the same form, but while *P* supports *C* in argument [Bismuth], *P* does not support *C* in argument [Wax]. This is a reformulation of Goodman's riddle (Goodman, 1983, §III.4). Norton concludes that, in trying to explain the differences between arguments [Bismuth] and [Wax], rule-based accounts must sacrifice universality for functionality; by adding extra conditions to their rules to block misapplications, the universality of the original rule is compromised (Norton, 2003, §2).

Norton's MTI, on the other hand, openly claims to be local. *P* supports *C* in argument [Bismuth] because bismuth is an element, and the argument is warranted by a local fact about elements, i.e., that chemical elements are *generally*[8] uniform in their physical properties. Argument [Wax], on the contrary, is not warranted by any local fact, given that "wax" is the generic name for various mixtures of hydrocarbons which are not uniform in their physical properties. By examining the local facts regarding the domain of each argument, claims Norton, we can understand whether it is warranted or not, and how.

Two important points about Norton's writings require clarification. First, Norton's MTI should be understood as a theory of inductive *warrant*, not as a theory of inductive support. While he is not often explicit about this distinction, this is an interpretation he acknowledges, for example, in (Norton, 2021a, p. 115). His central thesis is that the warrant of induction is wholly found in material facts. However, I (and many others) are concerned with *inductive support* itself. In other words; Norton is trying to understand what the *source* of inductive support is, whereas I am trying to provide an account of inductive support itself (in which, for the purposes of this chapter, I will adopt Norton's view on inductive warrant). If we read Norton as trying to provide an account of inductive support, we will find his MTI lacking, as I will point out in the next section; if we read it for what it is, we will find his MTI to be quite appropriate.

Second, Norton is concerned with the logic of induction, not with its epistemology. Some unclarity in Norton's writings, however, has resulted in several readers interpreting the MTI as an epistemic position. This confusion has obscured the value of the MTI. The key issue is that, although Norton's MTI is a view about the logic of induction, he does not always distinguish the logical

[8] As Norton points out, this "generally" is meant to capture the fact that some elements, such as sulphur, have different allotropic forms with different melting points.

from the epistemic matters clearly (as noted, for instance, by de Grefte (2020)). I believe this confusion is also caused, in part, by his usage of the term "inference", which he explicitly defines in a logical sense (Norton, 2021c, p. 20) but many readers interpret in an epistemic sense.

This confusion has resulted in an important line of criticism to Norton's MTI which can be settled here. Some critics have argued that Norton's MTI is not a clear/successful epistemology of induction (e.g. Davey, 2021; Livengood & Korman, 2020). Indeed, it is not, precisely because Norton's MTI is not a view on epistemology at all. The MTI is a view on inductive logic and, in particular, on the *warrant* of inductive arguments. The following fragment from Norton, in response to Davey (2021), is quite illustrative:

> [A] misunderstanding in Davey's text needs to be corrected. It labors over the problem of whether the material theory of induction is (or assumes) an internalist or an externalist epistemology; and settles on an externalist epistemology. The correct answer is "neither." The distinction between internalist and externalist epistemologies applies to accounts of the mode of justification of beliefs held by some agent. The material theory of induction concerns relations of inductive support among propositions, independently of whether these propositions are held as beliefs by some agent. (Norton, 2021a, p. 117)

Working out the epistemic consequences of the MTI (or of the HTI, for that matter) requires further work, similar to the work that de Grefte (2020) has urged and started to do. However, most of this work remains undone. This is not to say that this will be an unproblematic process, but it is only the products of this process that can be criticised from an epistemic standpoint, not the MTI itself. In conclusion, either the MTI is criticised from a logical standpoint, or some of its epistemic consequences are developed and then criticised from an epistemic standpoint – criticising the MTI from an epistemic standpoint rests on a misunderstanding of what the MTI is about.

We should now have a better understanding of Norton's MTI and the mechanism of factual warrant. In the next two sections I will argue that we need both rules (§1.5) and facts (§1.6) to do some work in a successful account of inductive support. I will then articulate this idea in the form of the HTI (§1.7).

1.5 Why We Need Rules

In this section I argue that rules of induction are necessary in order to have a *functional* model of inductive support. The key role for rules is to capture relations of inductive support. However, all rules of induction fail under certain circumstances, so no single rule results in a universal inductive logic. Therefore, no rule of induction is, by itself, an *accurate* model of inductive support.

The main claim of this section is that, without rules of induction, there is no account of inductive support. Rules of induction specify the conditions that an argument must fulfil so that the relation of support between P and C arises in that argument. Maybe it arises when P is an instance of C, or when C entails and best explains P, or when P raises the probability of C. However that may be, in specifying these conditions we are effectively providing a rule of induction that describes how *P* and *C* are connected. Every rule of induction, therefore, describes a particular mode of support that can arise between two propositions (e.g. between a piece of evidence and a hypothesis). When properly developed, these rules of induction become first-order models of induction, or inductive logics.[9]

We can contrast this with what Norton's MTI tells us about induction, since Norton's account fully relies on facts to do the work. The MTI tells us that *P* supports *C* in an inductive argument when, and *because*, the right facts obtain. However, the MTI never tells us what the right facts are (I will argue in §1.7 that we need rules in order to do that), hence, it is not a functional account of inductive support. Nevertheless, this question is outside the scope of the MTI. In other words; the MTI is not a functional account of inductive support because it is not an account of inductive support at all. The MTI is concerned with inductive *warrant* (Norton, 2021a, p. 115), and establishes that this warrant is factual, that is all (and it's a lot). Peden (2019) criticises the MTI for not providing an account of inductive support, but this is not a criticism anymore once we realise that we shouldn't expect the MTI to provide such an account. As I will argue in §1.7, if we want to integrate Norton's lessons on inductive warrant into an account of inductive support we need something like the HTI.

Of course, rules of induction may (and do) get it wrong sometimes – none of them results in a universal logic of induction. As Norton points out, for example, both arguments [Bismuth] and [Wax] follow the rule of enumerative induction, but *P* supports *C* only in argument [Bismuth]. Similarly, Bayesianism fails as a logic of induction under certain circumstances.[10] This point is developed in the next section. Hence, no single rule of induction is an *accurate* model of inductive support, although most rules are *locally* accurate in the right contexts.

In conclusion, we need to embrace rules of induction if we want a *functional* model of inductive support. In fact, it is precisely through rules that we can capture and describe relations of inductive support. However, no single rule of induction provides an *accurate* model of inductive support. The next section explores how factual warrant can help us here.

1.6 Why We Need Facts

In this section I argue that factual warrant is necessary in order to have an *accurate* model of inductive support. Thus, facts play a warranting role in our understanding of inductive support. In particular, we need facts in order to map the domains where specific rules of induction are accurate, that is, to define the domains where our rules of induction are properly capturing the relations of inductive support. In the previous section I have already argued that we need rules in order to characterize the relations of inductive support. Thus, I will here present the role of local facts *in addition to* the role of rules, not *instead* of them.

The first step in our defence of factual warrant is to realize that no rule of induction is universal, which is the main motivation behind Norton's MTI. Let's summarise some classical criticisms to

[9] In chapter 2, §2.3, I explain what it means for a rule to be properly developed.
[10] See (Norton, 2021c, ch. 10–16) for an extensive analysis of these failures.

the universality of rule-based theories of induction.[11] First is the family of "inductive generalization". Rule-based theories belonging to this family rely on the principle that an instance confirms the generalization. Theories in this family have a limited reach of evidence – i.e., there are many inductive arguments in which P supports C that are not captured by these theories, because the evidence is not an instance of the hypothesis (e.g. argument [Snowshoes]) . In addition, there are also many inductive arguments in which P does not support C despite following this principle (e.g. argument [Wax]), like Goodman famously noted.

Second is the family of "hypothetical induction". Rule-based theories belonging to this family rely on the principle that a hypothesis's ability to entail the evidence is a mark of truth. Nevertheless, this unrestricted principle allows for indiscriminate confirmation that we would not be willing to endorse. For instance, the hypothesis "all ravens are black and all swans are white" would be confirmed by the observation of a white swan, which is an undesired result since we would be confirming a hypothesis that bears information about the color of ravens by observing the color of swans. One notable way to "tame the indiscriminateness" of this principle is provided by versions of hypothetical induction known as *abduction* or Inference to the Best Explanation (IBE) (Aliseda, 2006; Lipton, 2004; Magnani, 2009). According to IBE, E supports H if H entails E and H is also the best explanation for E. IBE, of course, faces its own challenges, the most central of which may be to provide an account of what constitutes a *best explanation* that captures all relations of support accurately.[12] The prospects of getting IBE to be a universal account of induction are dim. Even Lipton, one of the leading defenders of IBE, acknowledged that "Inference to the Best Explanation cannot be the whole story about inference: at most, it can be an illuminating chapter" (Lipton, 2004, p. 4). As we will see, the HTI will provide the right framework in which to situate this "chapter".

Third, and last, is the family of "probabilistic induction". Rule-based theories belonging to this family rely on the principle that relations of support are probabilistic. It is unclear, however, whether relations of inductive support can always be described by our chosen probabilistic theory of induction. Sober, for instance, argues that Bayesianism fails when the values of probabilities cannot be empirically defended (Sober, 2008, p. 32), and he claims that we should then adopt another probabilistic theory of induction: likelihoodism. Norton provides a good summary of the diverse challenges that Bayesianism faces (Norton, 2021c, ch. 10–16). He argues, for example, that in some situations no probabilistic theory of induction is adequate, like in situations of completely neutral support (Norton, 2021c, pp. 348–53) or when our arguments are about some specific indeterministic systems like Norton's dome (Norton, 2010, §2). Challenges to the universality of probabilistic theories are abundant and diverse, and we don't have any reasons to expect that any single probabilistic theory of induction can successfully respond to all the challenges.

Hence, I echo Lipton's conclusion: rules of induction are "both too permissive and too strict, finding inductive support where there is none and overlooking cases of genuine support" (Lipton, 2004, pp. 17–18). In other words, rules of induction have "problems of scope", and thus they are not accurate theories of induction. While these problems can be addressed by restricting

[11] These are popular critiques and are well developed, for instance, in (Earman & Salmon, 1992), (Lipton, 2004) or (Norton, 2021c).

[12] See Lipton (2004) or Norton (2021c, Ch. 8–9) for some of these challenges and attempts at addressing them. Other authors have argued that IBE is not a fundamental argument form (Khalifa, Millson, & Risjord, 2017) or that "explanatoriness is evidentially irrelevant" (Roche & Sober, 2013).

the domain of validity of each rule of induction, this cannot be done without an appeal to factual warrant. In defining the domain where a given rule accurately captures the relations of inductive support, we are effectively providing a set of conditions that must obtain for our rule to be appropriate: thus, we are exposing the warranting fact of this rule.

One could still argue that not *all* rules of induction face problems of scope, but only those that have been articulated and examined so far. Furthermore, since the classification of rule-based theories offered here is not exhaustive, there are certainly some rules that have been left unaddressed. The one true universal rule of induction – would go this rebuttal – is still to be found. This wouldn't be a particularly strong counter-argument, but it can be put to rest here by realising that there is a fundamental reason why all rules of induction must fail as described above:[13]

> *Any logic of induction must restrict what happens in ways that go beyond logical consistency*. Hence, a logic of induction is applicable in some domain if the facts of that domain match the factual restrictions of the logic of induction. Since there is no universally applicable factual restriction, in general, different domains require different inductive logics. (Norton, 2021c, p. 337, emphasis in original)

In conclusion, we need to embrace factual warrant if we want our model of inductive support to be *accurate*. In the next section I go into more detail about how rules and facts come together in the HTI.

1.7 A Hybrid Theory of Induction

In this section I introduce the basics of the HTI. The HTI is a hybrid second-order model of inductive support. It is a *hybrid* model of inductive support in the sense that, in it, two ingredients play a necessary role in understanding inductive support: rules and facts. It is a *second-order* model of inductive support because it is a model within which first-order models of inductive support (i.e. logics of induction) can fit. In other words, the HTI is not itself directly involved in modelling inductive support, which is a job left for logics of induction. In what follows I will use the terms "rules of induction", "logics of induction" or "first-order models of induction" interchangeably. The main thesis of the HTI is the following: rules of induction accurately describe relations of inductive support when they are warranted, and facts do the warranting work.

Let's illustrate the HTI. I will focus on objective Bayesianism, since it is a relatively uncontroversial example. Most philosophers, even opponents to Bayesianism as a universal theory of induction like Norton or Reiss, will grant that Bayesianism works well when the outcome spaces and probabilities of our system are well-defined (or, as Sober puts it, probabilities are "empirically defensible"). Savage refers to these systems as "small worlds" (Savage, 1972), and the rest of systems are labelled "large worlds". A toss of a fair coin is a paradigmatic example of a small world, in which the outcome space is $\{H, T\}$ (H for "heads", T for "tails") and the probability of each outcome is 0.5. In the terms of the HTI, then, let us *assume* that the warranting fact for Bayesianism is the fact that our system is a small world. Thus, the domain of Bayesianism is limited to small worlds

[13] In this fragment Norton summarises an argument he develops in more detail in (Norton, 2021c, ch. 2)

and, correspondingly, Bayesianism accurately describes the relations of inductive support in that domain.[14]

Therefore first, and crucially, before understanding an inductive argument in terms of Bayesianism, we need to know whether our system is a small or a large world. This is not always an easy task. Norton's dome (Norton, 2008) is a highly idealized indeterministic system consisting of a point mass at perfect rest in the apex of a perfectly symmetrical frictionless dome. Newtonian mechanics tells us that the point mass can stay at rest in the apex forever or spontaneously move at any instant; we know nothing else.

Norton and Earman disagree about how we should characterize this system. According to Norton (2010, §2), this system cannot be described in probabilistic terms, since the physics of the system do not provide physical chances for its possible futures. Norton concludes that Bayesians, thus, cannot analyse this system. Earman (2020), however, disagrees, and suggests that quantum mechanics offers a solution for Norton's dome; furthermore, Earman then defends a version of Bayesianism as an account of induction in quantum mechanics.

This is a debate about whether the dome has well defined outcome space and probabilities and, therefore, it is a debate about whether Norton's dome is a small world or a large world – it is this feature about our system that determines whether Bayesianism accurately captures the relations of support in arguments about this system.

One might worry, at this point, that if even Bayesianism fails sometimes and, therefore, there is no universal logic of induction, then there cannot be a single understanding of inductive warrant either. However, this is not right. Inductive warrant can have a single nature, but be instantiated differently in different domains. This is what I have argued for in this chapter. In other words, the source of inductive support (the warrant) is always of the same *type*, i.e., factual; however, *which* particular facts are doing the warranting work will be different in different domains. For systems that are large worlds, inductive logics other than Bayesianism may be adequate. Those logics will require their own warranting facts. Norton (2021c, Ch. 16) argues, for example, that for certain types of quantum systems a particular inductive logic other than Bayesianism is required. The logic provided is a quantum inductive logic, it is not probabilistic, and it requires that our system be a particular type of quantum system – this is its warranting fact.

We can now see that the HTI has enough tools and structure to be an accurate and functional model of inductive support. First, the HTI tells us that any given rule will only provide accurate information about inductive support when it is warranted, that is, when the right facts about our system of interest obtain. This allows us to restrict the domain of application of a given rule, thus solving its problems of scope. Second, the HTI tells us that within the right domain, a rule of induction accurately describes the relations of inductive support. In this way, the functionality of rules is maintained. Peden (2019, p. 678), for instance, provides a list of questions about inductive support that Bayesianism can tackle and the MTI cannot. As I have already argued, when properly interpreted, the MTI is not expected to tackle these issues, since it is not a theory of inductive support. However, once we integrate Norton's MTI into the HTI, we realise that all the tools of Bayesianism are still available to us, albeit only in the right domain, as they should. In conclu-

[14] This is, of course, a simplification. Chapter 4 explores Bayesian inductive logics from the perspective of the HTI in much more detail.

sion, the HTI can address the main challenges of both rule-based and non-rule-based accounts of induction.

Within the HTI, rules of induction have the further role of guiding us to the necessary warranting facts, guidance that the MTI is lacking (as noted, for example, by Genta (2020, p. 21)). Let's recall Norton's analysis of arguments [Bismuth] and [Wax]. According to Norton, P supports C in argument [Bismuth] because the right warranting fact obtains, which is not the case for argument [Wax]. But how can we determine whether the warranting fact for an argument obtains if we don't know what we are looking for? This problem arises from thinking of *individual* arguments as being warranted by facts, an interpretation that Norton often endorses: "treat each inductive inference as a unique individual, each with its own special properties, as opposed to homogeneous instances of a single argument form" (Norton, 2011, p. 25). The HTI, however, urges us to think of *rules* being warranted by facts, not individual arguments. The warranting fact of a rule is the condition (or set of conditions) under which that rule is true. Exposing these conditions may not be easy, but we know what we are trying to expose and we must only do it once for every rule of induction.

Let's see how we would go about exposing the warranting fact for the rule of enumerative induction.[15] The rule of enumerative induction, as understood by Norton (2005), tells us that "an instance confirms the generalization". The HTI asks us to expose the conditions under which this rule is accurate; when is it true that an instance confirms the generalization? Here is a plausible answer: this rule is true when "the property being generalized is approximately uniform across the scope of the generalization". In the context of argument [Bismuth], it is true that "[c]hemical elements are *generally* uniform in their physical properties" (Norton, 2003, p. 651, emphasis in original), and therefore, it is true that "the property being generalized is approximately uniform across the scope of the generalization". However, this is not true in the context of argument [Wax]: the property being generalized is *not* approximately uniform across the scope of the generalization, since "wax" is the generic name for various mixtures of hydrocarbons which differ in their physical properties. Hence, argument [Bismuth] is warranted while argument [Wax] is not. While Norton states the same conclusion, the MTI does not offer a way to arrive at it and the HTI does.

We can repeat this analysis for any rule of induction. I have already offered a tentative analysis of Bayesianism in terms of the HTI above. If Sober is right, then Bayesianism is an accurate account of inductive support whenever probabilities are empirically defensible, i.e., whenever our system is a "small world". In that case, our system being a small world is the warranting fact for Bayesianism, and this is what we need to check before understanding an argument in terms of Bayesianism; this is the guidance that the HTI provides. Norton and Earman's debate about the dome is a debate about whether this warranting fact obtains for that specific system.

Here is a further example on the type of guidance that the HTI provides in our search for warrant. Some arguments in the context of RCTs are understood as following the rule that "correlation is a sign of causation". We can see appeals to this rule in classic accounts of causal inference from RCTs, like that of Papineau (1994) or Cartwright (2007). This rule tells us that evidence that "X is correlated with Y" supports the conclusion that "X is a cause of Y". However, the rule that "correlation is a sign of causation" is not universal – it is not always appropriate. There are many

[15] The analysis I provide here is only preliminary, but still useful for illustrative purposes. In chapter 2 I provide a detailed analysis of enumerative induction in terms of the HTI.

spurious correlations that are not a sign of causation. In asking when this rule is appropriate, we are effectively trying to expose its warranting fact. The search for the conditions under which correlation is indeed a sign of causation has been going on for a long time, and it remains a controversial topic. Classic voices in this debate (e.g. Byar et al., 1976; Papineau, 1994; Schwartz, Flamant, & Lellouch, 1980) argue that this rule is appropriate whenever a balance assumption obtains, i.e., the assumption that "all confounding causes are balanced between both arms of the trial".[16] I believe, however, that Fuller (2019) and Larroulet Philippi (2022) offer a more refined version of the warranting fact needed here. For example, Larroulet Philippi suggests that this rule is appropriate whenever the *total contribution* of confounders is approximately balanced between both arms of the trial (Larroulet Philippi, 2022, pp. 157–58), and balance in every confounder is not required. These articles by Fuller and Larroulet Philippi are perfect examples of the kind of work we must do in order to expose the warranting fact of a rule of induction. The HTI then tells us the following: we can use whatever warranting fact we end up exposing to check whether an argument from a claim of correlation to a claim of causation is warranted. If we agree with Larroulet Philippi, for any of these arguments we need to check whether the total contribution of confounders is approximately balanced between both arms of the trial. This may be a very hard thing to check, but the HTI tells us what warranting fact we need, something that the MTI cannot tell us by itself.[17]

In conclusion, the HTI articulates the role of both rules and facts in understanding and capturing relations of inductive support. Rules of induction describe the relations of support in the right domain, and guide us to the necessary warranting facts for a given argument; facts provide warrant for rules of induction, so they define the domain of applicability of rules. In this way, we obtain a functional and accurate model of inductive support.

1.8 The HTI in the Literature

In this section I show that the HTI is not a completely new idea, but an articulation of a view that has long been present in the literature but has not been made explicit. I do so by showing that many influential philosophers within the rule-based tradition have tried to expose the factual warrant for their rules of choice, and in doing so have adopted a view that can be very naturally understood in terms of the HTI. I finish by showing how the HTI is also implicit in Norton's writings. Hence, hopefully the HTI articulates a view that is shared by most philosophers working on induction and can serve as a common framework to situate debates more clearly.

Let's start by looking at Mill's analysis of enumerative induction. This is a classic example of an author looking for the *factual* warrant of a rule of induction. Although Mill states that he is concerned with "induction",

> Mill identifies induction with generalisation from experience, and specifically with the following sort of case: we have observed several members of a class, and found that they all have some property in common, with no exceptions. We then wish to infer that all members of that class have the property in question. (Graves, 1974, pp.

[16] Confounding causes are causes of the study outcome Y other than the intervention X.

[17] Again, this is only a preliminary analysis; see chapter 3 for a detailed analysis of a rule of causal inference in the context of RCTs.

303–304)

Therefore, in my terms, it is only enumerative induction that Mill is concerned with. Mill tries to warrant the rule of enumerative induction by appealing to the principle of the uniformity of nature (Mill, 1843/1974, 3.3.1).[18] Roughly, this principle states that "nature is uniform".[19] However, nature is not uniform, so Mill failed to show that enumerative induction is a universal rule of induction. Still, what matters for us here is that Mill looks for the warrant of his rule of interest in the outside world, not in other rules or a priori principles. Mill tries to determine how *the world* should be so that the rule of enumerative induction is valid, and this is precisely the perspective on induction that the HTI is capturing.

Quine's suggestion to solve Goodman's riddle is another case in point. Goodman notoriously pointed out that enumerative induction was an appropriate rule in some circumstances (e.g. argument [Bismuth]) but not in others (e.g. argument [Wax]). The response to this challenge, according to Quine (1969), is that enumerative induction is a valid rule when the right terms in the arguments are natural kinds. Roughly, to say that a term is a natural kind is to say that "it corresponds to a grouping that reflects the structure of the natural world" (Bird and Tobin, 2022). I ignore here all the problems associated with natural kinds in general and with Quine's response in particular. Again, what matters for us here is that Quine sought to warrant a particular rule of induction by exposing the conditions on *the world* and its structure such that his rule of interest worked. Quine's approach also fits the HTI's understanding of induction.

The work of many contemporary authors can also be understood as an effort towards exposing the warranting facts of their rule of interest. As introduced above, Norton (2010), Reiss (2020) or Sober (2008), argue that Bayesianism is only applicable when probabilities can be empirically defended, i.e., when we are in a "small world".

I don't pretend to argue here that any of these authors have exposed the right warranting facts for their rules of interest, but only to point out that this kind of work is commonplace, and it is precisely the kind of work that the HTI predicts and urges us to do. In doing this work, these authors accept that rules of induction *sometimes* describe relations of inductive support accurately – in particular, when the right facts obtain. These are the two central commitments of the HTI.

I also believe that the HTI is implicit in Norton's writings. As previously discussed, Norton's MTI is a view about inductive *warrant*, not about inductive support. His central thesis is that the warrant of induction is wholly found in material facts, with which I agree. However, there are a few fragments in his writings that go beyond the MTI and seem to suggest something like the HTI, e.g.:

> [A] logic of induction is applicable in some domain if the facts of that domain match the factual restrictions of the logic of induction. (Norton, 2021c, p. 337)
>
> I also welcome formal theorists who decide to justify their schemas with material facts. In so doing, they have become material theorists. That is not a defeat, but victory for the material theory. (Norton, 2021a, p. 124)

[18] Cited by book.part.section.

[19] This principle, of course, has been challenged on many grounds, mostly on the grounds that is is too vague. I will not be engaged in this debate here, since it is merely Mill's attempt that I need for my argument.

> The justificatory chain may pass through intermediate formal schemas of limited scope. (Norton, 2021a, p. 115)
>
> If used at all, the schema would have a purely intermediate role. [...] That is, a cascade of warrants may pass through a schema. The cascade terminates in facts that are the final warrant of the inference. (Norton, 2021c, p. 64)

In these fragments Norton talks about "inductive logics" or "schemas" being warranted by facts, as opposed to individual inductive inferences being warranted by facts, which is the language he uses most often. However, he never develops this view, but merely states it. Furthermore, he remains silent with respect to what role, if any, is left for rules of induction in understanding inductive support. This silence is fine, since Norton is providing an account of inductive *warrant*, and he explicitly says that rules do not play a warranting role (Norton, 2021c, p. 64), so that is all he needs to say. Still, Norton seems friendly to a view like the HTI.

However, while Norton admits that the justificatory chain (or cascade of warrants) *may* pass through a schema or rule, the HTI tells us that it *must* if we are to have a functional account of inductive support. As I have argued above, we need rules of induction in order to capture relations of inductive support, and they also guide us to the necessary warranting facts for a given argument.

In short, the HTI is not a completely novel idea but an articulation of a view that is widespread in the literature. In particular, the HTI is also implicit in Norton's writings, although never developed or defended. In this chapter I have taken this widespread view, articulated it clearly and provided some reasons to accept it. This is a necessary first step so that we can develop this view further.

1.9 Advantages of the HTI

In this section I summarise a few key advantages of the HTI. Some of its advantages have already been developed above: unlike rule-based and factual accounts of induction, the HTI provides a functional and accurate account of inductive support; unlike the MTI, the HTI provides guidance to the necessary warranting facts for a given argument. Let's examine some further advantages of the HTI.

By bringing into focus the role that rules play in our understanding of inductive support, the HTI can address some criticisms against the MTI vindicating the need for rules in a theory of induction (e.g. Baker, 2020; Genta, 2020). Baker, for example, raises the following concern: "Why this insistence that a theory of induction either be completely formal or completely material? I am not sure. In many cases, it seems natural to appeal both to formal and to material facts in order to explain the merits of a given inductive inference" (Baker, 2020, p. 118). This is a fair concern.[20] The HTI tells us that, indeed, we need to appeal to both rules and material facts to explain the merits

[20]To be clear, though, I disagree with the conclusion that Baker reaches immediately after raising this concern: "Sometimes the relevant facts are only formal. But it's hard to see how the assessment of an inductive inference could ever depend solely on material facts. It is because of this asymmetry that my sympathies lie with the formal side in this debate" (Baker, 2020, p. 118). I believe the solution is to think about inductive support as the HTI specifies, which is compatible with Norton's adequate response to Baker (Norton, 2021a, pp. 114–15).

of an argument. In fact, we cannot understand inductive support if we don't look at inductive arguments from the perspective of a warranted rule of induction, as argued in §1.5. For example, we might say that *P* supports *C* in argument [Bismuth], in the first place, because *P* is an instance of *C* and the rule of enumerative induction tells us that an instance confirms the generalization. This appeal to a warranted rule is necessary in explaining the merits of this argument. However, the fact that *P* is an instance of *C* is relevant here only because the rule of enumerative induction is warranted in this domain. The warranting work is done by a fact, the fact that "generally, chemical elements are uniform in their physical properties", and that is the source of support. The HTI describes the ways in which rules and facts are necessary in order to understand the relation of inductive support in an argument.

The HTI can also be useful in tackling some problems in epistemology. For instance, Lipton distinguishes two central questions about our principles of induction: the question of description and the question of justification (Lipton, 2004, p. 7). The question of description asks what principles we use in making inductive inferences; the question of justification asks whether these are good principles. He then focuses on the question of description. After examining some potential principles of induction in the literature, including Bayesianism, Lipton concludes that all these approaches are "finding inductive support where there is none and overlooking cases of genuine support" (Lipton, 2004, pp. 17–18). Thus, understood as principles of inductive inference, they are not good descriptions of our inferential practices. Lipton's diagnosis is that none of these approaches gives "enough structure to the black box of our inductive principles to determine the inferences and judgements we actually make" (Lipton, 2004, p. 18). I believe the HTI provides the necessary missing structure. A preliminary response to this challenge would be the following: we use inductive principles (rules of induction) that are warranted in the domain of our inference. The black box of our inductive principles cannot contain just one universal inductive principle, or else we will find that this black box has problems of scope, as Lipton concludes. Instead, by including the mechanism of factual warrant we can delimit the domain of validity of each principle. Lipton's black box requires rules (a plurality of them) and warranting facts for those rules; this is its internal structure. In this way, we may be able to provide a more accurate description of our inductive practices. I believe the HTI can help in this and other debates in epistemology, but exploring this in detail is a matter for future work.

The HTI can also help us understand the bigger picture of induction, i.e., how different rule-based theories of induction can fit together. From the perspective of the HTI we can understand how different rule-based theories relate to each other by mapping the domains where they are valid. Some rule-based theories might complement each other nicely: Sober suggests, for example, that Bayesianism works where probabilities are empirically defensible, and when they're not, we should be likelihoodists (Sober, 2008). Or maybe we will find out that one theory of induction can be reduced to another, which is applicable in the same domain. Several authors (e.g. Day & Kincaid, 1994; Henderson, 2014; Iranzo, 2008) have explored the possibility of reducing IBE to Bayesianism. Day and Kincaid, in particular, argue that IBE can be understood in terms of a more fundamental theory of induction (maybe Bayesianism, they suggest), such that IBE is valid within a subdomain of the domain of this more fundamental theory.[21]

[21] Day and Kincaid argue that "[...] IBE does not name a fundamental pattern of inference but that it is instead an instance of another, more general inference strategy. That strategy infers to warranted beliefs from background information and data. When those background beliefs essentially involve claims about explanation and those claims

The HTI also fits very well with the idea that there is no unique scientific method but a plurality of them. This idea has been present at least since Suppes suggested it back in his 1978 presidential address to the Philosophy of Science Association (Suppes, 1978). Many philosophers have embraced one form or another of scientific pluralism since then. If there are a plurality of scientific methods (Kellert, Longino, & Waters, 2006; Suppes, 1978), for example, it should not come as a surprise that there may be a plurality of inductive logics. Some philosophers of science have taken a metaphysical turn, defending that there is some sort of pluralism in reality itself. Cartwright (1999), for instance, argues that we live in a "dappled" world that we should understand as a patchwork of laws. In our world, correlations and causal interactions are local, and thus scientific theories have a local scope. If the world is "dappled", we have good reasons to expect both a plurality of scientific methods and, relatedly, a plurality of inductive logics. The HTI, however, does not *commit* us to any of these pluralisms – it just fits nicely with them.

The HTI can also help us clarify existing debates about inductive logic by bringing competing views under a common framework. For example, Norton takes himself to be arguing against all rule-based theorists. According to Norton, the MTI "differs fundamentally from virtually all approaches to inductive inference in the present literature" (Norton, 2021c, p. v). On the one hand, Norton argues that inductive arguments are warranted by material facts. On the other hand, most rule-based theories of induction make no explicit mention to factual warrant.

This divide, I argue, can and should be dissolved. Recall that Norton concedes the following to rule-based theorists: "I also welcome formal theorists who decide to justify their schemas with material facts. In so doing, they have become material theorists" (Norton, 2021a, p. 124). Rule-based theorists will also acknowledge, even if implicitly, the importance of factual warrant in their accounts. This can already be seen in Mill and Quine's attempts at addressing the limitations of enumerative induction; similarly, it is commonplace to assume that certain causal arguments from RCTs require a balance assumption. These assumptions are all warranting facts, necessary for the application of the corresponding rules.

We should thus understand the disagreement between Norton and rule-based theorists as a difference in focus. Norton is focusing on one necessary ingredient in our understanding of inductive support (the facts), while rule-based theorists are focusing on another (the rules). Each ingredient has a role in a successful account of inductive support. I believe both parties can agree on something like the HTI, and realise how they can benefit from the work of each other. In short, the HTI brings all these views into a common framework and can therefore help us tackle disagreements with more clarity. Many disagreements will persist, of course, but some may be resolved.

1.10 Conclusion

In this chapter I have introduced the basics of the HTI. This theory is hybrid in the sense that, in it, two ingredients play a necessary role in our understanding of inductive support: rules and facts.

---

ground inferences from the data to new beliefs, this general inference strategy becomes IBE. What exactly this more general inference strategy involves is, of course, a matter of debate. It may, for example, be given a Bayesian reading or be understood as a kind of bootstrapping" (Day and Kincaid, 1994, p. 282)

According to rule-based theories of induction, inductive arguments are warranted by rules. Rules of induction are necessary in order to have a functional model of inductive support. However, all rules of induction fail under certain circumstances, so no single rule provides a universal inductive logic. Hence, no rule of induction is, by itself, an accurate model of inductive support. Rule-based theories cannot define their domains of validity without appealing to factual warrant.

According to non-rule-based theories of induction, inductive arguments are not warranted by rules; Norton, in particular, argues that inductive arguments are warranted by material facts. Factual warrant is an accurate warranting mechanism, since it is local. However, we still need rules in our account of induction if we want it to be functional, since rules are responsible for describing the relations of inductive support.

I have argued that both rules and facts play a necessary role in a successful account of inductive support, and I call the resulting account a Hybrid Theory of Induction. According to the HTI, rules of induction accurately describe relations of inductive support when they are warranted, and a rule of induction is warranted if the right facts about the matter of the induction obtain. I don't take the HTI to be a completely novel idea, but rather an articulation of a widespread view on induction that has not been explicitly articulated yet. Unlike rule-based and factual accounts of induction, the HTI provides a functional and accurate account of inductive support. Unlike the MTI, the HTI provides guidance to the necessary warranting facts for a given argument. The HTI can also help us tackle some problems in epistemology, and it provides a useful general framework to examine and understand induction and to make sense of how different rules of induction can coexist. Moreover, the HTI allows us to clarify and resolve some current debates on induction, like the apparent divide between Norton and rule-based theorists.

For the purposes of this chapter, I have embraced Norton's MTI as an account of factual warrant. However, Norton's MTI faces some challenges. I believe an alternative account of factual warrant, according to which warrant comes in degrees, should be devised and integrated into the HTI. I turn to this task in the next chapter.

Chapter 2

Logical Virtues of a Graded Account of Factual Warrant

2.1 Introduction

In chapter 1 I argued that a functional and accurate account of inductive support needs two ingredients: rules and facts. The main purpose of chapter 1 was to show that both of these ingredients are necessary. The main purpose of this chapter is to refine that account.

Rule-based accounts of induction have dominated the literature for many years, so their underlying rules have been explored and discussed in detail. Therefore, I here take up the more neglected task of developing the role of facts in understanding inductive support. Although the idea that matters of fact may play a role in understanding inductive support is not new, Norton's MTI is the only contemporary well-developed account of factual warrant available. While, for the purposes of chapter 1, I embraced Norton's MTI, in this chapter I develop a different Graded account of Factual Warrant (GFW). The crucial difference between Norton's MTI and the GFW is that, while Norton's MTI is a binary account of warrant, according to the GFW factual warrant comes in degrees. In this chapter I develop this account of warrant and integrate it into the HTI.

This chapter has three goals. First, and foremost, to introduce the GFW and the resulting version of the HTI. Second, to provide two arguments in favour of the GFW and the updated HTI. This will be accomplished by highlighting two logical virtues of the GFW (§2.7 and §2.8). Later on, chapter 3 will highlight some epistemic virtues of these accounts. The third and last goal is to illustrate the GFW and the resulting version of the HTI. This will be accomplished by analysing the rule of enumerative induction in terms of the HTI. This rule and the resulting analysis will be used as a running example throughout this chapter. Enumerative induction is a simple and popular rule, and, despite its limitations, it is widely used in both everyday and scientific reasoning. Therefore, it will be interesting to tackle this rule of induction first. In chapters 3 and 4 I address other rules of induction.

Here is how I will proceed. In section 2.2 I briefly summarise the key ideas of this chapter. In the rest of the chapter I flesh out these ideas more precisely. In section 2.3 I articulate my understanding of rules of induction as sentence schemas, which is necessary in order to articulate the GFW. In section 2.4 I introduce the notion of "verisimilitude", which will be important in articulating the GFW. In section 2.5 I introduce the GFW and integrate it in the HTI. In section 2.6 I argue that, while inductive support is interest-dependent, the HTI is interest-independent. In section 2.7 I argue that the HTI, by embracing the GFW, illuminates the connection between an argument's factual warrant and the strength of that argument, which is an advantage with respect to the MTI. In section 2.8 I focus on a criticism from Reiss to Norton regarding the role of idealisations (and theory) in induction; I then show how the GFW can address Reiss' criticism more clearly than Norton's MTI. Section 2.10 concludes.

2.2 Summarising the Arguments

The arguments in this chapter will be introduced with the help of a few formalisms. Before doing so, in this section I summarise the key ideas of this chapter in a more intuitive manner.

Let's recall argument [Bismuth] first:

P : Some samples of the element bismuth melt at 271ºC
C : All samples of the element bismuth melt at 271ºC [Bismuth]

This argument can be very naturally understood in terms of the rule of enumerative induction, which tells us, roughly, that if "some As are B, then all As are B".[1] Thus, enumerative induction, understood in this way, tells us that the degree of support in any argument from "some As are B" to "all As are B" (like argument [Bismuth]) is maximal or, in other words, that *P* entails *C* in such arguments.

However, this is not always the case. The rule of enumerative induction is not always accurate. In chapter 1 I argued that rules of induction accurately describe relations of inductive support when they are warranted, and a rule of induction is warranted if the right facts obtain. Here, I refine this account; I argue that, for each argument, there is a given set of *ideal* conditions under which the information on inductive support that a rule provides is perfectly accurate. For argument [Bismuth] and the rule of enumerative induction, the ideal conditions are the following: "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature". Under these ideal conditions, enumerative induction is perfectly accurate with respect to argument [Bismuth], since the support in that argument is indeed maximal (as the rule says).

Furthermore, not all cases in which these ideal conditions fail to obtain are the same: we can be closer or further from meeting them and, accordingly, our rule of interest will be more or less accurate. In the case of argument [Bismuth], when "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature" then the degree of support is maximal and our rule is fully accurate. We move away from these ideal conditions as less samples of

[1]There are, of course, other ways in which one can articulate the rule of enumerative induction. For present purposes this simple form will suffice. Later in this chapter I articulate and discuss the rule of enumerative induction in more detail.

bismuth are like the observed samples with respect to their melting temperature.[2] As this happens, the degree of support in argument [Bismuth] decreases and, therefore, it becomes less accurate to say that the degree of support in argument [Bismuth] is maximal, which is what our rule says. Thus, as we move away from the ideal conditions for enumerative induction in argument [Bismuth], our rule becomes less accurate. This is the core idea behind the GFW.

I introduce the GFW in §2.5; I then highlight two logical virtues of this view in §2.7 and §2.8. In §2.7 I answer the following question: which features of the world are relevant for the degree of support of a given argument? The GFW offers an answer: those features of the world that determine how close we are to meeting the ideal conditions for our rule and argument of interest. For the example above, the ideal conditions tell us that "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature". The relevant features of the world in this case are, for instance, the degree of uniformity in the physical properties of chemical elements, or the existence of allotropes of bismuth which might differ in their physical properties. This intuitive answer does not follow from other understandings of warrant, but it follows from the GFW. In §2.8 I point out that the GFW offers a clear picture of the role of idealisations (and theory) in our understanding of inductive support. The MTI has some trouble accommodating the role of idealisations, and I echo a discussion between Reiss and Norton to illustrate this. In contrast, according to the GFW idealisations and theory can (and often do) play the role of "ideal conditions". This is a natural and unproblematic understanding of their role in our account of inductive support.

In this section I have briefly summarised some of the central ideas of the chapter, avoiding the use of formalisms. I will now articulate these ideas more precisely. In order to do so, in the next section I introduce my understanding of rules of induction as sentence schemas.

2.3 Rules as Sentence Schemas

The only commitment I made in chapter 1 with respect to the nature of rules of induction is that they describe relations of inductive support. However, it is now important to be a bit more specific about what rules of induction are. In this section I lay out my understanding of rules of induction: rules of induction are sentence schemas whose instances inform us about the relations of inductive support in an argument.

Following the logical tradition, I will understand rules of induction as schemas. A *schema* is a linguistic template together with a side condition for using it (Corcoran & Hamid, 2016). The linguistic template contains placeholders and significant words; the side condition specifies how to fill out these placeholders and it may also specify how the significant words or symbols are to be understood. We obtain an *instance* of a schema when the placeholders in the linguistic template are filled out according to the side condition. We classify schemas based on the type of instances they produce.[3]

[2] This is a simplification. What it means exactly to move away from certain ideal conditions is dependent on our interests. I expand on this in §2.6.

[3] My understanding of rules of induction as sentence schemas (i.e., linguistic templates that once filled out produce sentences) might remind some readers of a logical positivist project. However, this is not an accurate assessment, and I encourage these readers to continue reading. During this section and the rest of the chapter it will become evident that

Rules of induction are either conceptualized as sentence schemas or as argument-text schemas. Sentence schemas are schemas the instances of which are sentences; argument-text schemas are schemas the instances of which are argument-texts. An argument-text is a system of sentences, one of which we call the conclusion and the rest of which we call premises. An argument is that which is expressed by an argument-text, just like a proposition is that which is expressed by a sentence. Following is the well known argument-text schema corresponding to *modus ponens*, a rule of deductive logic:

P_1 : If A then B
P_2 : A [Modus Ponens Argument-Text Schema]
C : B

with the side condition that A and B are declarative sentences of English.

Current literature is rarely explicit about the nature of rules of induction, and rules are often conceptualized as argument-text schemas. In this way, the rule of enumerative induction is commonly understood as the argument-text schema

P : Some As are B
C : All As are B

with the side condition that As are entities and B is a property of those entities.

The use of argument-text schemas in conceptualizing rules of logic can be traced back, at least, to Aristotle; Aristotle's mood BARBARA, for example, referred to the argument-text schema

P_1 : P belongs-to-every M
P_2 : M belongs-to-every S
C : P belongs-to-every S

with its corresponding side condition (Corcoran & Hamid, 2016, §4). But BARBARA is a rule of deductive logic, and such rules can adequately be understood as argument-text schemas, since the relations of support in deductive arguments depend purely on their syntax.

However, problems arise when we try to capture rules of induction as argument-text schemas; the solution is to understand rules of induction as sentence schemas. While this understanding might be implicit in some contemporary texts, I have not come across an explicit articulation of rules of induction in these terms. In the remaining part of this section I briefly develop and justify why we should understand rules of induction as sentence schemas.

First and foremost, it is important to realise that any rule which can be captured by an argument-text schema can also be captured by a sentence schema. This can be done by generating a sentence schema corresponding to the conjunction of the premises of the argument-text schema of interest, followed by the appropriate word representing the support conferred to the conclusion (e.g. the word "therefore" to represent implication) and finally followed by the conclusion of the

syntactic considerations play no role in my understanding of inductive support.

argument-text schema. In this way, the rule of *modus ponens*, which is captured by [Modus Ponens Argument-Text Schema], can also be represented as the sentence schema "If A then B, and A, therefore B" with the same side condition. Thus, by choosing to understand rules of induction as sentence schemas instead of argument-text schemas we are not losing any generality.

However, sentence schemas can capture any rule of induction, while argument-text schemas cannot. In particular, sentence schemas, but not argument-text schemas, allow us to capture rules of induction that do not impose any restriction on the content of the premises and conclusions of the arguments they apply to. Bayesian inductive logics are the most notable of such rules.

Let's contrast Bayesianism with enumerative induction. Enumerative induction does impose some restrictions on the type of information that must be present on the premises and conclusions of the arguments it applies to. For such arguments, the premise must express information about some property of a particular subset of entities of a certain type, and the conclusion must generalize this property to all entities of that type. Consequently, enumerative induction could be articulated both as a sentence schema as an argument-text schema.

During this chapter I will use the rule of enumerative induction as a running example, so I shall now articulate it first in the form of a sentence schema; I will then articulate it as an argument-text schema. This rule is often taken to be the sentence schema "some As are B, therefore all As are B", with the side condition that As are the entities over which we are generalizing and B is the property being generalized. However, let's be a bit more precise in our articulation of this rule; I will take the rule of enumerative induction to be the following schema,

EIS: if all Xs in S are such that $p(X) = Y$, then all Xs are such that $p(X) = Y$

with the side condition that Xs are the entities over which our argument of interest is generalizing, S is the subset of Xs from which our argument is generalizing, p is the property of Xs being generalized, and Y is the value of p for all Xs in S.

I call this rule EIS because it is the rule of Enumerative Induction under a Strong interpretation. This is a strong interpretation of enumerative induction because, according to EIS, all Xs in set S being such that $p(X) = Y$ *implies* that all Xs are such that $p(X) = Y$, while weaker interpretations of the rule do not talk about *implication* but about *confirmation*.[4]

As stated above, since EIS does impose restrictions on the content of the premises and conclusions of the arguments it applies to, it can also be articulated as an argument-text schema:

$$\begin{array}{ll} P: & \text{All } X\text{s in } S \text{ are such that } p(X) = Y \\ C: & \text{All } X\text{s are such that } p(X) = Y \end{array} \qquad \text{[EIS Argument-Text Schema]}$$

with the same side condition as the corresponding sentence schema EIS I introduced above.

Thus, arguments [Bismuth], [Bismuth*] and [Salt] can all be understood in terms of the argument-text schema [EIS Argument-Text Schema].

[4] I favour this strong interpretation of the rule and the resulting analysis, but this analysis can be repeated for any other interpretation. In chapter 3, §3.2, I provide a "strong" and a "weak" interpretation of a rule of induction, and their corresponding analyses.

P :	Some samples of the element bismuth melt at 271ºC	[Bismuth]
C :	All samples of the element bismuth melt at 271ºC	

P :	The melting temperature for some samples of bismuth is 271ºC	[Bismuth*]
C :	The melting temperature of bismuth is 271ºC	

P :	This sample of radium chloride is monoclinic	[Salt]
C :	All samples of radium chloride are monoclinic	

This is because, for each of these arguments, there is one way to fill in the placeholders in [EIS Argument-Text Schema] such that the resulting argument-text will express our argument of interest. It is important to differentiate arguments from argument-texts again. An argument is that which is expressed by an argument-text, just like a proposition is that which is expressed by a sentence. Thus, [Bismuth] and [Bismuth*] are different argument-texts but they express the same argument. This is because the propositional content of their premises and conclusion coincides. [Bismuth] and [Salt] are different argument-texts expressing different arguments. All argument-texts which can be understood in terms of enumerative induction are such that the arguments they express can be generated by the schema [EIS Argument-Text Schema].

However, not all rules will allow for this; for example, it is not possible to articulate Bayesianism in the form of an argument-text schema, because it is not the case that all argument-texts which can be understood in terms of Bayesianism are such that their arguments can be generated by any single argument-text schema. Compare argument [Bismuth] with the following argument:

P :	This card has been drawn at random from a standard 52-card deck and it is red	[Card]
C :	This card is a heart	

Arguments [Bismuth] and [Card] can both be understood in terms of Bayesianism, but there is no single argument-text schema that can generate both of these arguments. The challenge here is that a diversity of arguments, whose premises and conclusions have nothing in common, can be understood in terms of Bayesianism. However, Bayesian inductive logics can be articulated in the form of sentence schemas – this is a project I tackle in chapter 4.

Two remarks are important at this point; first, the problem in understanding rules of induction as argument-text schemas is not due to the fact that the syntax of inductive arguments does not determine their degree of support. The arguments expressed by argument-texts with different syntax can actually be understood in terms of the same argument-text schema, as long as these arguments can be generated by the same argument-text schema. In this way, arguments [Salt] and [Bismuth*], despite being expressed by argument-texts which differ in syntax, can both be understood in terms of the same argument-text schema [EIS Argument-Text Schema], as I have argued above. The problem with argument-text schemas is that they do not allow us to capture rules of induction that place no restrictions on the content of the premises and conclusions of the arguments they apply to. Without these restrictions, there can be no single argument-text schema that is able to generate all the arguments for which this rule is applicable. Bayesianism and IBE are notable examples of such rules.

The second remark is that, as we should expect, we can find rules of induction being conceptualized as sentence schemas in the literature. Norton (2005), for instance, expresses hypothetico-deductivism as the principle that "the ability to entail the evidence is a mark of truth", which is shorthand for the sentence schema "the ability of H to entail E is a mark of the truth of H" with the side condition that E and H are declarative sentences of English and express, respectively, the premise and the conclusion of our argument of interest. Accordingly, IBE, as a development of hypothetico-deductivism, is often understood as the sentence schema that "the ability of H to entail E and best explain it is a mark of the truth of H", with the corresponding side condition.

In chapter 1 I referred to rules of induction as describing relations of inductive support; however, I shall now refine this statement. Strictly speaking, since rules of induction are not sentences but sentence *schemas*, they cannot inform us about anything until their placeholders are filled out. It is the *instances* of rules which inform us about inductive support.

Let us look at rule EIS, and explore how we would go about generating instances of this rule. Rule EIS has the following placeholders: X, S, p and Y. Once these placeholders have been filled in according to the side condition, we obtain an instance of rule EIS. Thus, we could think of EIS as a function of four variables – X, S, p and Y – and we could refer to instances of EIS as EIS(X, S, p, Y). However, we should realise that all of these placeholders are filled in with information about our argument of interest, as specified in the side condition of rule EIS. Thus, we can also think of EIS as a function of a single variable: our argument of interest A. This will be more convenient. Hence, I will refer to instances of EIS as EIS(A).

However, rules of induction can also contain placeholders which are not filled in with information about our argument of interest – see, for example, the rules I explore in chapter 4.[5] Thus, in general terms, I call R(A, $\Phi_1, \ldots, \Phi_n$) the instance of rule R evaluated in context (A, $\Phi_1, \ldots, \Phi_n$), where A refers to our argument of interest, and $\Phi_1, \ldots, \Phi_n$ are all placeholders in R which are *not* filled in with information about our argument of interest.

Let us see what an instance of rule EIS looks like, then. EIS([Bismuth]), for example, is the result of filling in the placeholders in schema EIS in the context of argument [Bismuth]. In this case, X are samples of bismuth, S is the set of observed samples of bismuth, p is the melting temperature, and Y is 271ºC. Thus, EIS([Bismuth]) is the sentence "if all samples of bismuth in the set containing the observed samples of bismuth are such that their melting temperature is 271ºC, then all samples of bismuth are such that their melting temperature is 271ºC", or, equivalently, "if the observed samples of bismuth melt at 271ºC, then all samples of bismuth melt at 271ºC".

As mentioned above, it is instances of rules that inform us about the inductive support in an argument. EIS([Bismuth]) is informing us about the relation of support in argument [Bismuth]; in particular, it tells us that P implies C in this argument, so the degree of support is maximal. However, as established in chapter 1, the information that the instances of a rule provide is not always accurate – rules get it wrong sometimes. EIS([Wax]), for instance, tells us that "if the observed samples of wax melt at 91ºC, then all samples of wax melt at 91ºC", which is obviously wrong

[5]The Bayesian rules I explore in chapter 4 contain a placeholder, $\mathcal{M}$, which must be filled in with a probability model. Such probability models are not contained in our arguments of interest – they must be independently provided.

since wax is not uniform in its physical properties. Thus, while EIS([Bismuth]) seems very plausible, EIS([Wax]) doesn't. In the next section I introduce the GFW, which will help us deal with this challenge.

Last, it is important to note that an argument A can only be understood in terms of those rules for which $R(A, \Phi_1, \ldots, \Phi_n)$ can be formed. If sentence $R(A, \Phi_1, \ldots, \Phi_n)$ cannot be formed, rule R cannot inform us about the relation of inductive support in argument A. Argument [Bismuth] can be understood in terms of EIS because EIS([Bismuth]) can be formed, as developed above. However, argument [Card] cannot be understood in terms of EIS, since EIS([Card]) cannot be formed; placeholder X, for example, cannot be filled in because there is no such thing as "the entities over which our argument of interest is generalizing" in the context of argument [Card], which is not generalizing anything. Instead, argument [Card] can be understood in terms of other rules, like those of Bayesianism. This just means that enumerative induction (EIS) is not even a potential mode of support for argument [Card], which is quite an intuitive result since argument [Card] is not generalizing any property.

In this section I have introduced my understanding of rules of induction. Rules of induction are sentence schemas, the instances of which are sentences which inform us about relations of inductive support. An argument A can only be understood in terms of rule R if $R(A, \Phi_1, \ldots, \Phi_n)$ can be formed. As established in chapter 1, the information about inductive support that the instances of rules provide is not always accurate. In section 2.5 I introduce the GFW, which will help us deal with this. But, first, I must introduce the notion of verisimilitude, which will be important in articulating the GFW. In the next section I turn to this task.

2.4 Verisimilitude

In this section I introduce some basic commitments regarding the notion of "verisimilitude" (or "truthlikeness"). This notion will be important in articulating the GFW and, therefore, in articulating the HTI. The notion of verisimilitude is needed, for example, in order to capture the intuition that some false propositions are closer to the truth than other false propositions. The false proposition "there are 9 planets in the solar system", for instance, is closer to the truth than the false proposition "there are 900 planets in the solar system". Hence, we can say, the former proposition is more truthlike than the latter. I will take the standard contemporary approach to verisimilitude, and conceptualize it as a real-valued function v, such that proposition p_1 is closer to the truth than proposition p_2 if and only if $v(p_1) > v(p_2)$ (Vassend, 2020, p. 1361).

However, the exact form that a measure of verisimilitude v will take in a given context is dependent on our interests. This feature of verisimilitude is widely recognized in the literature (Miller, 1994; Niiniluoto, 1998; Northcott, 2013; Vassend, 2020). It is easy to see why verisimilitude is interest-dependent. The verisimilitude of a proposition p, $v(p)$, is a measure of the closeness to the truth of p or, in other words, of the degree of similarity between p and the truth (Oddie, 2016, §1.1). However, as Niiniluoto points out, the degree of similarity between p and the truth "depends on the respects of comparison that are taken to be relevant" (Niiniluoto, 1998, p. 14), and on the weight we assign to the dissimilarity in each respect. Take a simple example: the proposition that "the earth is a sphere". It is a well-established fact in physical geodesy that the earth's polar axis is shorter than its equatorial diameter by a fraction of about 1/298. The precise value of this flattening is of great scientific relevance: it explains two perturbations in the lunar orbit, periodic changes to the orientation of the earth's axis, and – since it jointly results from the earth's rotational inertia and the mutual attraction of its constituent particles – it offered a primary test for Newtonian gravitational theory (Ohnesorge, 2022). Hence, for scientists interested in the motion or gravitational potential of the earth, minute differences between the length of its polar axis and its diameter are very relevant. Using Northcott's terminology (Northcott, 2013), for these scientists the "seriousness of error" in assuming that the earth is a sphere is huge, since their practices would be heavily impacted by the errors in this assumption.[6] Accordingly, for these scientists the proposition that "the earth is a sphere" is *not* very truthlike. For most other purposes, including medium-distance navigation or mapping entire countries, the actual polar flattening of the earth is irrelevant. After all, even Isaac Newton noted that "the inequality [in the length] of degrees [at different latitudes] is so small that in geographical matters the shape of the earth can be considered to be spherical", just after he had predicted that the earth is flattened at the poles (Newton, 1723/1999, Book 3, §20). For a cartographer, then, the "seriousness of error" in assuming that the earth is a sphere is very small, since her practices will not be impacted by the errors in this assumption. Hence, given the interests and purposes of a cartographer, the proposition that "the earth is a sphere" is very truthlike.[7]

This interest-relativity of verisimilitude can be seen as a special case of a more general phenomenon: the interest-relativity of any measure of similarity. First, we must note that measures of verisimilitude are measures of similarity of a specific type. In particular, a measure of the verisimilitude of p is a measure of the similarity between p and the truth. Second, we must note that any measure of similarity is interest-relative. This is because any two entities can be compared with respect to multiple properties, each of which may be differently important. Is a small green stapler more similar to a small red stapler or to a big green stapler? We cannot answer this question without specifying which properties of a stapler we care about, and how much we care about them. Accordingly, "[s]imilarity judgements [...] require not only the selection of relevant properties but also the weighting of their importance" (Morreau, 2010, p. 470).

Admittedly, we are still lacking a successful account of the concept of truthlikeness and its properties.[8] Popper's work (1963) initiated two closely related approaches to understanding the nature of truthlikeness: the *content* approach and the *consequence* approach. In the content approach (later developed, for example, by Miller (1978) or Kuipers (1987a)) the truthlikeness of a proposition depends on its amount of true and false content. In the consequence approach (later developed, for example, by Burger and Heidema (2005) or Gemes (2007)) the truthlikeness of a proposition depends on its amount of true and false consequences. Besides facing other challenges, these approaches, in the original Popperian form, had the terrible consequence that a false proposition could never be closer to the truth than another false proposition (Miller, 1974; Tichý, 1974). In response to this challenge, Tichý (1974) and Hilpinen (1976) argued for a *likeness* approach to truthlikeness. According to this view, a proposition's truthlikeness depends on the similarity be-

[6]Northcott tells us that "interest-relativity might be seen merely as an instance of the seriousness of errors problem, broadly construed. In any case, both imply that once a theory has fallen short of full truth, thereafter there just is no univocal answer as to how *much* it has fallen short" (Northcott, 2013, p. 1482).

[7]I thank Miguel Ohnesorge for useful discussion about this example.

[8]See (Oddie, 2016) or (Kuipers, 1987b) for a summary of the main challenges related to truthlikeness.

tween the possible worlds that make the proposition true and the actual world. However, likeness accounts are in need for an appropriate likeness function on possible worlds. Furthermore, Miller (1974, 1976, 2006) has notably criticised likeness accounts on the grounds that they are language-dependent – that is, that we obtain different measures of verisimilitude when we translate our statements (or scientific theories) to different languages. Thus, there is no consensus with respect to the nature of truthlikeness or how to best articulate this notion.

Thus, I take verisimilitude to have the following features: it is a real-valued function measuring the similarity between a proposition and the truth, it is interest-relative, and it is a contested notion. However, the HTI does not require a detailed and complete account of verisimilitude. Whatever the nature of verisimilitude may be, and however we choose to measure it, the HTI only requires us to accept the following commitments: (1) that some propositions are closer to the truth than other propositions and (2) that there is some way to measure the truthlikeness of a proposition. The first of these commitments is the fundamental intuition behind the notion of verisimilitude. In the second commitment I am following the bulk of the literature on verisimilitude (see (Niiniluoto, 1998) for a good overview of such proposed measures). Much more work is required in order to advance our understanding of verisimilitude, but this work is not a pre-requisite for the HTI.

2.5 A Graded Account of Factual Warrant

In this section I articulate the GFW, which is an account of inductive warrant, and I integrate it within the HTI, which is an account of inductive support. I lay out these views at the outset for convenience, since I will refer to them throughout this chapter, but I do not defend them in this section. My defence of the GFW and the resulting version of the HTI spans chapters 2 and 3. In this chapter, sections 2.7 and 2.8 will each provide a different argument in favor of the GFW and the resulting version of the HTI. Both arguments provided in this chapter point out logical virtues of the GFW. In the next chapter, instead, I point out some epistemic virtues of the GFW.

The GFW is an account of inductive warrant, that is, of the source of inductive support. Since, in the GFW, the warranting work is done by facts, it is an account of *factual* warrant, just like the MTI. However, while the MTI is a binary account of warrant, the GFW is, as its name suggests, graded. Thus, according to the GFW, factual warrant comes in degrees. The MTI tells us that an argument is warranted if the right local fact obtains, and not warranted otherwise; the GFW tells us that there is no such thing as a "right" warranting fact for an argument, but that, instead, many facts can warrant an argument, each to a varying degree.

Let me introduce the concept of $\mathrm{IWC_R}$, the Ideal Warranting Conditions for rule R. $\mathrm{IWC_R}$ will allow me to articulate the GFW and the HTI precisely. Like R, $\mathrm{IWC_R}$ is also a sentence schema. I call $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ the instance of schema $\mathrm{IWC_R}$ evaluated in context $(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$. I define $\mathrm{IWC_R}$ as the sentence schema such that, for all arguments A such that $\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ can be formed, $v(\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)) = v(\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n))$. In this chapter I will expose $\mathrm{IWC_R}$ for the rule of enumerative induction; in chapter 3 I will expose $\mathrm{IWC_R}$ for a rule used for causal arguments in the context of RCTs; in chapter 4 I will expose $\mathrm{IWC_R}$ for Bayesian inductive logics.

I can now articulate the GFW and the HTI precisely, in terms of IWC_R:

> GFW: given a rule R there is an associated sentence schema IWC_R such that, given any argument A such that $R(A, \Phi_1, \ldots, \Phi_n)$ can be formed, $v(IWC_R(A, \Phi_1, \ldots, \Phi_n))$ is a measure of the degree to which rule R is warranted in context $(A, \Phi_1, \ldots, \Phi_n)$.
>
> HTI: given a rule R there is an associated sentence schema IWC_R such that, given any argument A such that $R(A, \Phi_1, \ldots, \Phi_n)$ can be formed, $v(R(A, \Phi_1, \ldots, \Phi_n)) = v(IWC_R(A, \Phi_1, \ldots, \Phi_n))$.

I have now introduced the general forms of the GFW and the HTI. However, since the rules I explore in this chapter and the next do not contain any placeholders $\Phi_1, \ldots, \Phi_n$ (which are *not* filled in with information about our argument of interest), I will omit all references to $\Phi_1, \ldots, \Phi_n$ for now, for convenience and brevity of notation. I will go back to the general forms of the GFW and the HTI in chapter 4, where it is required.

Let me now express the core idea of the GFW in more intuitive terms. Recall that the warrant of a rule is the source of inductive support. My account of inductive warrant is factual, which means that, in it, the warranting work is done by facts. Thus, the source of support of an argument is a particular fact. The main idea for my graded account of factual warrant is that the full source of support that our rule requires can be present "to a varying degree". $IWC_R(A)$ represents the "full" source of support for rule R in argument A, i.e., the conditions under which R(A) is true. The idea of partial warrant is that this source of support can be more or less present; $v(IWC_R(A))$ captures the degree to which this source of support, $IWC_R(A)$, is present.

Thus, the HTI provides a way to obtain information about the relation of inductive support in an argument. It tells us that, given a rule R, the information provided by R(A) is truthlike to degree $v(IWC_R(A))$. We can restate this contribution in epistemic terms. An agent interested in argument A does not know $v(R(A))$. Sentence R(A) describes the relation of support in argument A, and an agent interested in argument A does not know the degree of support in argument A, since this is precisely what is at stake. By providing a sentence other than R(A) whose verisimilitude is the same (i.e. $IWC_R(A)$), the HTI allows the agent to figure out $v(R(A))$. In other words; for every statement about inductive support that we may be interested in, the HTI provides another statement about *the world* whose verisimilitude is the same.

The HTI tells us that there is a schema IWC_R for each rule R. The kind of work required to expose IWC_R is commonplace, despite not being phrased in my terms, and I now include two examples of such work, which I already anticipated in chapter 1. The literature on RCTs has been grappling with this question for a while, exploring the conditions under which the rule that "correlation implies causation" is valid (e.g. Byar et al., 1976; Cartwright, 2007; Papineau, 1994). Roughly, the rule of interest for these authors can be expressed as the schema "if X is correlated with Y, then X caused Y" with the side condition that, in the context of an RCT, X is the intervention and Y is the outcome measured. Classic literature argues that IWC_R, in this case, is the schema that "no cause other than X has had an impact on Y" with the same side condition. This IWC_R is often called a "balance assumption". Similarly, as I pointed out in chapter 1, Quine argued that enumerative induction is a valid rule when the entities over which we are generalizing are natural kinds, and the property being generalized is a natural property. The rule of enumerative induction is often understood as the schema that "some As are B, therefore all As are B", with the side condition that As are the entities over which we are generalizing and B is the property being generalized. Quine's suggestion, then, is that IWC_R in this case is the schema that "As form a natural kind, and B is a natural property of As", with the side condition that As are the entities over which we are generalizing and B is the property being generalized. Thus, the work of these authors can be very naturally understood as an attempt at exposing IWC_R for their rule of interest. Unfortunately, there is no universal approach to exposing IWC_R for any given rule R; each rule will pose its own intellectual challenges. In the rest of this chapter, as well as in chapters 3 and 4, I expose IWC_R for several rules, providing detailed examples of how such an analysis looks like.

To finish the section, let me address a potential objection to the HTI that may arise at this point. Crucial to both the GFW and the HTI is the notion of IWC_R, which I have *defined* as the sentence schema such that, for all arguments A such that R(A) can be formed, $v(IWC_R(A)) = v(R(A))$. It might seem, then, that the HTI is merely restating this definition. However, this is not true. What the HTI is telling us is that, in fact, such an IWC_R (as defined) exists for any rule R. This is a significant claim. It tells us that we can, and should, think about induction in a very specific way. Of course, I cannot provide a conclusive defence of the HTI, whatever that would be, but I will advocate the HTI in several ways: chapter 1 has already listed several advantages of having both rules and facts play a role in our account of inductive support (§1.9); the current chapter points out further logical virtues of the – now fully fleshed-out – HTI; the next chapter points out some epistemic virtues of the HTI; last, chapters 2, 3 and 4 each contain an example of a well-known rule of induction analysed in terms of the HTI, illustrating that IWC_R for our rules of interest can be exposed, and that doing so is useful.

In this section I have introduced the GFW and the corresponding version of the HTI. I have not defended these accounts here; I will do this in the rest of this chapter and in chapter 3. In the next section I analyse the rule of enumerative induction in terms of the HTI. This will allow me to address an important concern about the role of verisimilitude in my account, while illustrating the type of analysis that the HTI allows us to do.

2.6 Interest-Independence of the HTI

In this section I show the following: the fact that the verisimilitude of a sentence is dependent on our interests is not problematic for the HTI. The HTI tells us that, given a rule R, there is an associated sentence schema IWC_R such that, given any argument A for which R(A) can be formed, $v(IWC_R(A)) = v(R(A))$. Some readers may be doubtful that an equality such as $v(IWC_R(A)) = v(R(A))$ can even be affirmed, since it is unclear how to establish the value of these verisimilitudes and, furthermore, verisimilitude is dependent on our interests (as I introduced in §2.4). In this section I show that, given a rule R, a schema IWC_R can be found such that this equality holds independently of our interests. I do this by exposing IWC_R for the rule of enumerative induction. This is obviously not a knockdown argument, but it should illustrate that schemas IWC_R as required by the HTI exist. In chapters 3 and 4 I illustrate this for two additional rules of induction. This section has a second aim, which is to illustrate the way in which we can analyse a rule of induction from the perspective of the HTI; I provide a detailed analysis of EIS, which will be useful during the rest of the chapter.

A preliminary clarification is required. What I will argue in this section is that the HTI is independent of our interests, not that *inductive support* is. In fact, since verisimilitude is interest-relative, inductive support will be interest-relative too. This is because the verisimilitude of the information on inductive support that R(A) provides, $v(\mathrm{R(A)})$, will be interest-relative. However, what is required for the HTI to be interest-independent is that the equality $v(\mathrm{IWC_R(A)}) = v(\mathrm{R(A)})$ holds regardless of our interests, even if $v(\mathrm{R(A)})$ is itself interest-relative, and this is what I will argue in this section.

2.6.1 Enumerative induction in terms of the HTI

Understanding a rule R from the perspective of the HTI means specifying the pair of sentence schemas R and $\mathrm{IWC_R}$. In this section I will provide an analysis of enumerative induction in terms of the HTI. I specified the sentence schema EIS above, so it is now time to expose $\mathrm{IWC_{EIS}}$.

Recall that EIS is the schema "if all Xs in S are such that $p(X) = Y$, then all Xs are such that $p(X) = Y$", with the corresponding side condition. In this section I will show that $\mathrm{IWC_{EIS}}$ is the schema "all Xs are like the Xs in set S with respect to $p(X)$", where all placeholders must be filled in the same way as in EIS. Let me highlight both of these schemas for convenience:

> EIS: if all Xs in S are such that $p(X) = Y$, then all Xs are such that $p(X) = Y$
>
> $\mathrm{IWC_{EIS}}$: all Xs are like the Xs in set S with respect to $p(X)$

with their corresponding side conditions.

Let's see why $\mathrm{IWC_{EIS}}$ is this particular schema. In order to argue that $\mathrm{IWC_{EIS}}$ is this schema I have to show that for all A such that EIS(A) can be formed, $v(\mathrm{IWC_R(A)}) = v(\mathrm{R(A)})$. In what follows, I will show that this equality holds, independently of our interests, for a fixed A. It will become apparent that the upcoming analysis is independent of the particular A chosen, as long as EIS(A) can be formed.

My argument of choice is argument [Bismuth]. EIS([Bismuth]) can be formed, and it is the sentence "if the observed samples of bismuth melt at 271ºC, then all samples of bismuth melt at 271ºC". Correspondingly, $\mathrm{IWC_{EIS}}$ ([Bismuth]) is the sentence "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature".

Before continuing, it will be useful to restate EIS([Bismuth]) and $\mathrm{IWC_{EIS}}$ ([Bismuth]) together for reference, since I will use these sentences quite often in what follows:

> EIS([Bismuth]): if the observed samples of bismuth melt at 271ºC, then all samples of bismuth melt at 271ºC
>
> $\mathrm{IWC_{EIS}}$([Bismuth]): all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature

If we have exposed the right $\mathrm{IWC_{EIS}}$, then $v(\mathrm{IWC_{EIS}}\ ([\mathrm{Bismuth}])) = v(\mathrm{EIS}([\mathrm{Bismuth}]))$. However, determining the verisimilitude of these sentences is not straightforward. Imagine a scenario, call it scenario 1, where every unobserved sample of bismuth melts at a different temperature within 10^{-30}ºC of 271ºC, but none melts exactly at 271ºC. Imagine another scenario, call it scenario 2, where all unobserved samples of bismuth melt exactly at 271ºC except for one, which melts at 5,000ºC. In both scenarios, the observed samples of bismuth melt at exactly 271ºC. Is IWC_{EIS} ([Bismuth]) more truthlike in scenario 1 or in scenario 2? The answer to this question, of course, will depend on our interests, as I advanced in §2.4. Still, if I have properly identified IWC_{EIS} then, whichever our interests are, once they are fixed, $v(IWC_{EIS}(A)) = v(EIS(A))$. Let's fix our interests and check whether this is so.

2.6.2 First set of interests

I will here specify a given set of interests and illustrate that, while inductive support is dependent on those interests, this is not problematic for the HTI since $v(IWC_{EIS}(A)) = v(EIS(A))$ still holds. In other words; once our interests are fixed, any good measure of $v(IWC_{EIS}(A))$ will also be a good measure of $v(EIS(A))$.

Let's look at $v(IWC_{EIS}$ ([Bismuth])). IWC_{EIS} ([Bismuth]) tells us that "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature". Let's now fix our interests. Suppose that all we care about is whether all samples of bismuth are *exactly* like the observed samples of bismuth with respect to their melting temperature. If an unobserved sample differs from one of the observed samples with respect to its melting temperature we do not care how big or small that difference is. In this case, an appropriate measure of $v(IWC_{EIS}$ ([Bismuth])), given our interests, might be

$$\frac{n_T}{N} \tag{2.1}$$

where T is the melting temperature of the observed samples of bismuth, n_T is the number of samples of bismuth that melt at T and N is the total number of samples of bismuth.[9] This is because, given our interests, we will consider IWC_{EIS} ([Bismuth]) to be more truthlike the more samples of bismuth that melt at T, independently of the properties of those samples that do not melt at T. IWC_{EIS} ([Bismuth]) will be maximally truthlike when all samples of bismuth melt at T and minimally truthlike when no sample of bismuth beyond the observed ones melts at T. The measure n_T/N behaves in the right way, taking a maximum value of 1 when all samples of bismuth melt at T and a minimum value of n_o/N (where n_o is the number of observed samples of bismuth) when no sample beyond those observed melts at T.

As desired, if n_T/N is an appropriate measure of $v(IWC_{EIS}$ ([Bismuth])) given our interests, it will also be an appropriate measure of v(EIS([Bismuth])) given the same interests. EIS([Bismuth]) tells us that "if the observed samples of bismuth melt at 271ºC, then all samples of bismuth melt at 271ºC". In the same way as before, this statement will be more truthlike the more samples of bismuth that melt at T, independently of the features of those samples that do not melt at T. Again, EIS([Bismuth]) will be maximally truthlike when all samples of bismuth melt at T and minimally truthlike when no sample of bismuth beyond the observed ones melts at T. The measure n_T/N behaves in the right way, taking a maximum value of 1 when all samples of bismuth melt at T and a minimum value of n_o/N when no sample beyond those observed melts at T.

[9] Of course, individuating samples of bismuth is not straightforward, but my point is independent of how we choose to do that.

We should also realise how inductive support is dependent on our interests. EIS([Bismuth]) tells us that P implies C in argument [Bismuth] or, equivalently, that inductive support is maximal in argument [Bismuth]. Still, v(EIS([Bismuth])) is dependent on our interests; for this set of interests, n_T/N is an appropriate measure of v(EIS([Bismuth])) and, therefore, it is an appropriate measure of the verisimilitude of the claim that inductive support is maximal in argument [Bismuth]. The larger v(EIS([Bismuth])), the more truthlike it is that inductive support is maximal in argument [Bismuth], and, therefore, the larger the degree of support in argument [Bismuth] actually is. Hence, how much P supports C in argument [Bismuth] depends on v(EIS([Bismuth])) and, therefore, it depends on what our interests are.

An alternative measure of verisimilitude, given the same interests, might be $(n_T - n_o)/(N - n_o)$. This measure would behave very similarly to n_T/N, with the difference that it would take a minimum value of 0 for the case in which no sample beyond those observed melts at T. If we find that this measure better captures our interests, then it is easy to check that it turns out to be a good measure of both $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) and v(EIS([Bismuth])).

To conclude: while v(EIS([Bismuth])) is dependent on our interests and, therefore, so is the degree of support in argument [Bismuth], the equality $v(\text{IWC}_{\text{EIS}}([\text{Bismuth}])) = v(\text{EIS}([\text{Bismuth}]))$ still holds.

2.6.3 Second set of interests

Now suppose that we have different interests. Let's assume that now we don't simply care about whether any given sample melts *exactly* at the same temperature as the observed samples; we care about *how close* the melting temperature of any given sample is to the melting temperature of the observed samples. I will now show that, given these interests, it is still the case that any good measure of $v(\text{IWC}_{\text{EIS}}(\text{A}))$ will also be a good measure of v(EIS(A)).

In this case, an appropriate measure of $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) might be

$$-\sum_{i=1}^{N} |t_i - T| \tag{2.2}$$

where t_i is the melting temperature of sample i, N is the total number of samples and T is the melting temperature of the observed samples. This measure will take a maximum value of 0 when all samples melt at T, and a negative value that will decrease as the sum of the differences in melting temperature, $|t_i - T|$, becomes greater.[10]

As desired, if expression 2.2 is an appropriate measure of $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) given our interests, it will also be an appropriate measure of v(EIS([Bismuth])) given the same interests. Sentence EIS([Bismuth]) tells us that "if the observed samples of bismuth melt at 271ºC, then all samples of

[10] Whether there is a *minimum* value for measure 2.2 is unclear. This depends on whether there is an absolute hot, i.e., a hottest temperature that matter can reach. This is unclear, and it remains a controversial topic among physicists. If there is an absolute hot, then t_i can only get so far from T, since t_i would be bounded between absolute zero and absolute hot. If there is no absolute hot, however, measure 2.2 can take any value in $(-\infty, 0]$.

bismuth melt at 271ºC". Given our newly defined interests, this statement will be more truthlike the larger the value of expression 2.2.

Of course, many alternative measures of verisimilitude can be developed given our new set of interests. For instance, if we want to maintain the convention that $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) takes value 1 when it is maximally truthlike, we can modify our measure to

$$1 - \sum_{i=1}^{N} |t_i - T| \tag{2.3}$$

Alternatively, we could use the Mean Squared Error (MSE) of estimator T as our measure of verisimilitude, with the sign changed to respect the convention that larger values of v correspond to more verisimilitude:

$$-\frac{\sum_{i=1}^{N} (t_i - T)^2}{N} \tag{2.4}$$

This would be an appropriate choice, for example, if we care about larger deviations from T much more than about smaller deviations, since we will judge the verisimilitude of IWC_{EIS} ([Bismuth]) and EIS([Bismuth]) accordingly.

This analysis can be repeated for any argument A such that EIS(A) can be formed, and we will realise that any measure of $v(\text{IWC}_{\text{EIS}}$ (A)) that we find appropriate given our interests will also be appropriate as a measure of v(EIS(A)).

For this second set of interests our conclusion is the same: while v(EIS([Bismuth])) is dependent on these interests and, therefore, so is the degree of support in argument [Bismuth], the equality $v(\text{IWC}_{\text{EIS}}$([Bismuth])) = v(EIS([Bismuth])) still holds.

In sum, although v(EIS([Bismuth])) (and, thus, the degree of support in argument [Bismuth]) is dependent on our interests, the equality $v(\text{IWC}_{\text{EIS}}$([Bismuth])) = v(EIS([Bismuth])) is not.

2.6.4 Concluding remarks

A brief note on epistemology is necessary here. By now, some readers may be worried that, in order to determine the values of the measures of verisimilitude described above, we would need to have access to all information about the melting temperatures of all samples of bismuth. Of course, this would be problematic since we do not have access to such information – this is precisely why we are interested in an argument like [Bismuth] to begin with. However, we do not need access to such complete information.

What is the case, and widely acknowledged, is that in order to make an inductive inference we need *some* background knowledge. The HTI tells us more; we need this knowledge in order to estimate the value of the appropriate measure of verisimilitude. We can have more, or less, of the relevant background knowledge, and accordingly our epistemic risks will be lower or higher. In order to estimate the value of n_T/N, for instance, knowledge about the uniformity of the physical properties of chemical elements, as well as knowledge about the existence of allotropes of bismuth, would be relevant. As I point out in the next section, it is a virtue of the HTI that it specifies *which* features about the world are relevant for the degree of support of an argument.[11]

In this section I have argued that the HTI is interest-independent. This is an important observation, since verisimilitude is not interest-independent. This could threaten the HTI, which tells us that given a rule R there is an associated sentence schema $\mathrm{IWC_R}$ such that, given any argument A for which R(A) can be formed, $v(\mathrm{IWC_R}(\mathrm{A})) = v(\mathrm{R}(\mathrm{A}))$. The interest-relativity of v might cast some doubts on whether an equality such as $v(\mathrm{IWC_R}(\mathrm{A})) = v(\mathrm{R}(\mathrm{A}))$ can even be affirmed. However, I have argued that given a rule R, there is a schema $\mathrm{IWC_R}$ such that any measure of $v(\mathrm{IWC_R}(\mathrm{A}))$ that we find appropriate given our interests will also be appropriate as a measure of $v(\mathrm{R}(\mathrm{A}))$ given the same interests. I have argued this by showing that such an $\mathrm{IWC_R}$ can be exposed for a particular rule, that of enumerative induction. Of course, this is not a knockdown argument. In the next two chapters I will repeat this analysis for two other rules of induction, showing that $\mathrm{IWC_R}$ can also be exposed for those rules and that the interest-independence of the HTI is maintained. Thus, I hope to show that schemas $\mathrm{IWC_R}$ that behave as the HTI requires can indeed be exposed for our rules of interest, and that doing so is useful.

2.7 Factual Warrant and Inductive Support

In this section I pose a challenge to Norton's MTI and show how the HTI can respond to it, thanks to the GFW. The HTI illuminates the connection between factual warrant and inductive support, something that Norton has trouble doing by relying in his MTI. To be clear, as I argued in chapter 1 (and is explicit in my statement of the HTI above) our rules also play a role in understanding inductive support, a role which Norton does not acknowledge. My argument here is that, furthermore, he doesn't offer an account of the link between factual warrant and inductive support either, and the HTI can help us here.

As introduced above, Norton's understanding of factual warrant is binary. According to Norton, an argument (or inference, in his terms) is either warranted or it is not. Remember that Norton uses the term "inference" to refer to what I call "argument". According to Norton, an argument is warranted if the right local facts obtain, and not warranted otherwise. He sometimes uses the dichotomy between good and bad inferences too: "[a]ccording to [the material theory of induction], what separates the good from the bad inductive inferences are background facts – the *matter* of the inference, as opposed to its *form*" (Norton, 2021c, p. 23, emphasis in original). Thus, Norton says, arguments are either warranted or unwarranted; they are either good or bad.

Hence, according to Norton, argument [Bismuth] is warranted because the right fact obtains, i.e., the fact that "[c]hemical elements are *generally* uniform in their physical properties" (Norton, 2003, p. 651, emphasis in original).

Still, with the tools of the MTI, Norton wants to capture the idea that *P* can support *C* in an inductive argument to varying degrees. As discussed in chapter 1, Norton's MTI is an account

[11] In chapter 3 I make the equivalent epistemic point; the HTI specifies which knowledge about the world is relevant for our assessment of inductive support.

of inductive warrant, not an account of inductive support; furthermore, it is a *binary* account of inductive warrant. His suggestion is that we can find out about the degree of support in an argument by investigating the argument's warranting fact: "[t]he warrant for an induction is a fact, and we assess and then control the inductive risk by exploring and developing that fact" (Norton, 2021c, p. 48). For instance, in the case of argument [Bismuth], Norton says that what matters is the word "generally", which in this context means something like "in so far as allotropy does not interfere" since "[s]ome elements, such as sulfur, have different allotropic forms with different melting points" (Norton, 2003, p. 651). The word "generally" might mean different things in different warranting facts. Recall argument [Salt]:

P :	This sample of radium chloride is monoclinic	[Salt]
C :	All samples of radium chloride are monoclinic	

According to Norton, argument [Salt] is warranted by the Weakened Haüy's Principle, the fact that "generally, each crystalline substance has a single crystallographic form". According to Norton, "[i]n this case, 'generally' means 'in so far as polymorphism does not interfere.' So the nature of the risk one takes in accepting the conclusion will differ with each context" (Norton, 2021c, pp. 43–44).[12] Norton's suggestion, then, is that the inductive support of an argument is to be determined by investigating the argument's warranting fact.

A problem with Norton's response is that he gives us no way to reach this conclusion; he provides no account of what it is about an argument's warranting fact that is relevant to the strength of that argument, and no guidance about how our investigation of an argument's warranting fact should go. He says, about IWC_{EIS} ([Bismuth]), that what matters is the meaning of the term "generally" which, in this case, means "in so far as allotropy does not interfere". Why is this what matters for the strength of the argument in this case? It is unclear. It might be an intuitive analysis, but it does not follow from any general account of inductive support. Norton, indeed, tells us that there is no such account.[13]

However, the HTI, relying on the GFW, provides a clearer account of the link between factual warrant and inductive support. First, the HTI urges us to expose IWC_{R} ([Bismuth]). In order to do that, we need to expose IWC_{R} for the rule R that we choose to analyse argument [Bismuth]. Thus, we need to specify our rule of choice. I choose to analyse argument [Bismuth] in terms of EIS. However, this argument could also be analysed in terms of other rules, like those of Bayesianism. This is not problematic. Let me first analyse argument [Bismuth] from the perspective of EIS; I will then briefly show why analysing it in terms of Bayesianism, for example, would lead us to the same conclusions.

Here is the general answer that the HTI provides to the objection I have raised against Norton: the features of the world that are relevant for the inductive strength of argument A are those which determine $v(\text{IWC}_{\text{R}}(\text{A}))$, given our interests.

Let's assume that, given our interests, measure 2.1 is an appropriate measure of $v(\text{IWC}_{\text{R}}(\text{A}))$; then, those features about the world that bear on n_T/N are the ones relevant for the degree of

[12]Polymorphism is the condition in which a solid chemical compound can exist in more than one crystalline form.

[13]This problem is exacerbated by the fact that we are also missing any guidance to determine what the warranting fact of an argument is, a problem that I discussed in chapter 1 and that the HTI addresses.

support of argument [Bismuth]. As Norton points out, allotropy is relevant for the strength of argument [Bismuth], but he provides no account of why this is so. The HTI tells us that allotropy is relevant because allotropes of bismuth might differ in their melting temperature, which has an impact on the value of n_T. In this way, the HTI provides a clear logical picture of what it is about the world that matters for the inductive strength of a given argument.

I can now articulate the link between factual warrant and inductive support more explicitly. EIS tells us that *P* implies *C* in argument [Bismuth], hence, that the degree of support is maximal. However, this is only true when IWC_{EIS} ([Bismuth]) is true. When IWC_{EIS} ([Bismuth]) is false, *P* does not imply *C*, but not all cases in which IWC_{EIS} ([Bismuth]) is false are the same. Assuming that n_T/N is a good measure of $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) given our interests, then, the larger n_T/N is, the larger $v(\text{IWC}_{\text{EIS}}$ ([Bismuth])) will be and, according to the HTI, the larger v(EIS([Bismuth])) will be. EIS([Bismuth]) tells us that *P* implies *C* in argument [Bismuth]. Thus, the larger v(EIS([Bismuth])), the more truthlike it will be that *P* implies *C*, therefore, the more *P* will support *C*. In sum, the larger n_T/N, the more *P* supports *C* in argument [Bismuth], and the HTI provides a clear picture of this relation. Allotropy is relevant for n_T/N and, thus, it is relevant for the degree of support of argument [Bismuth] given our set of interests. This is how the HTI connects factual warrant and inductive support.

Importantly, we would reach the same conclusion if we chose to understand argument [Bismuth] from the perspective of a different rule. Take Bayesianism, for example. I have not yet offered a full account of Bayesianism in terms of the HTI (I do this in chapter 4) but, for present purposes, it suffices to realise that part of our rule (which I call BAY for Bayesianism) will be a probability model $\mathcal{M}$ of our system of interest, and IWC_{BAY} will be the commitment that model $\mathcal{M}$ is correct or adequate-for-purpose. If model $\mathcal{M}$ assumes that bismuth has no allotropes, it will tell us that *P* implies *C* in argument [Bismuth], just like EIS; however, the adequacy-for-purpose of this model (and, hence, the truthlikeness of IWC_{BAY} in our context) will depend on how much allotropy actually interferes. Alternatively, suppose that model $\mathcal{M}$ perfectly accounts for all allotropes of bismuth. Our rule will then be fully warranted, but the effects of allotropy in the uniformity of the melting temperature of bismuth will have already been accounted for within our rule, in our probability model. Allotropy will still impact the degree of support in argument [Bismuth], either by being accounted for in model $\mathcal{M}$ or, otherwise, by impacting the truthlikeness of IWC_{BAY} in our context.

In this section I have shown how the HTI illuminates the connection between factual warrant and inductive support. For a given set of interests, the facts about the world will determine the value of $v(\text{IWC}_{\text{R}}$ (A)). The HTI tells us that $v(\text{IWC}_{\text{R}}$ (A)) $=$ v(R(A)). R(A) provides information about the relation of inductive support in argument A and, therefore, $v(\text{IWC}_{\text{R}}$ (A)) determines the truthlikeness of this information about inductive support. Thus, thanks to the GFW, the HTI provides a clear logical picture of the way in which certain features of the world bear on the degree of support of a given argument.

2.8 A Response to Reiss

Reiss (2020) has recently raised an interesting objection to Norton's MTI. Reiss embraces Norton's negative thesis that there are no universally applicable rules of induction, and he also embraces Norton's positive thesis that material facts warrant inductive arguments. However, Reiss argues, material facts are not the *only* warrant of induction; there are at least six other types of warrant. In this section I acknowledge that Reiss raises a valid concern regarding the role of idealisations (and theory) in induction. I then show how the GFW can address Reiss' criticism more clearly than Norton's MTI.

Reiss argues that theories, idealisations, adequacy-for-purpose, ethical norms, methodological norms and conceptual norms also warrant inductive arguments.[14] I take Norton's response to Reiss' appeal to norms (ethical, methodological and conceptual) to be satisfactory (Norton, 2021a, pp. 122–23). The debate around the role of adequacy-for-purpose deserves greater attention, but this is a matter for future work. For now, I will focus on Reiss' appeal to the warranting role of idealisations and theory.

Following Reiss, I will take theories and idealisations to be hypotheses (or sets of hypotheses). According to Reiss, "[t]heories are bodies of substantive *hypotheses* used to systematize and unify a range of diverse phenomena. [...] Idealisations are more specific *hypotheses* that conflict with known facts (or are presumed or suspected to do so) but that are useful nevertheless" (Reiss, 2020, p. 12, my emphasis). Since I am addressing Reiss' objection, during this section I will adopt his understanding of theories and idealisations. The relevant feature about theories and idealisations, for the coming debate, is that, as hypotheses, they can be (or are known to be) false. I will make my point by focusing on a discussion between Norton and Reiss about the role of one specific idealisation, but my arguments are applicable to theories (understood as bodies of hypotheses) as well.

Reiss' criticism regarding the role of idealisations and theory as drivers starts by noting an ambiguity in Norton's writings: sometimes Norton acknowledges that theories or idealisations can warrant arguments, while most of the time he is explicit that only facts can play this role. However, facts and hypotheses are very different kinds of things, since facts must be true, but not hypotheses. Reiss is right in noting this ambiguity. The central tenet of Norton's MTI is that inductive inferences are warranted by *facts* which are, by definition, true. In case this is not clear, Norton is often explicit that only truths can warrant inferences: "[t]he material theory of induction arises from the recognition that the truth of these background factual presumptions is all that is needed for the inductive inference to be warranted" (Norton, 2021c, p. 64). Still, theories and idealisations can also play a role in our understanding of inductive support, and Reiss urges that the role they play is a *warranting* role, just like Norton seems to suggest in some passages. Here is one such passage:

[14] Let me clarify some terminology. Reiss uses the terms "warrant", "drive", "enable" and "justify" interchangeably. In addition, Reiss and Norton use the term "inference" instead of "argument". Norton is clear that by "inference" he means a relation of support between two propositions and, hence, this relation is for logic to study, nor for epistemology. Thus, he means by "inference" what I mean by "argument". I keep these terms distinct to separate logical from epistemic matters. Reiss is not explicit in his usage of the term "inference" but, since he is engaged in a direct response to Norton, I will assume here that he uses the term "inference" in the same sense as Norton.

> Cosmology seeks to discover the structure of the universe on the largest scale. If the universe is infinite in spatial extent, then the finite portion observable by us is minuscule. What we see is infinitely outweighed by what we cannot see. The essential assumption that allows us to proceed from what we can see to what we cannot is the "cosmological principle." It asserts that the universe is roughly homogenous in its large-scale properties. While this wording may seem somewhat vague, standard applications of the principle employ it unambiguously. In our vicinity of the universe, matter is distributed roughly uniformly in galaxies in a space of constant, possibly zero, curvature. The cosmological principle authorizes us to infer that this condition obtains everywhere in the whole universe. Much of modern cosmological theory proceeds from this authorization. (Norton, 2021c, p. 80)

However, the cosmological principle is an idealisation and, as such, it is false. Thus, there is a tension in Norton's writings, which usually tell us that only truths can warrant inferences, but sometimes tell us that idealisations can do that too.

In response to Reiss' criticism Norton has conceded the following:

> I do allow idealizations to serve the role of a warrant. Reiss recalls one: the cosmological principle, which, construed narrowly, asserts the falsehood that the universe is exactly homogeneous and isotropic. Idealizations such as these are quite admissible as warrants, in so far as the falsities in them do no compromise the inductive inference to be warranted. (Norton, 2021a, p. 122)

We can raise several questions about Norton's response. We can, for example, wonder about the meaning of the following claim: "the falsities in [the idealisations] do not compromise the inductive inference to be warranted". Norton proceeds to provide a couple of examples to illustrate this claim. However, this issue is not directly related to the topic at hand and providing a good account of these examples would require considerable space. Hence, let me push aside this concern for now.

For present purposes I will focus on another aspect of Norton's response: the aforementioned tension between facts and idealisations in the MTI. In his response to Reiss Norton writes the following: "I do allow idealisations to play the role of a warrant" (Norton, 2021a, p. 122). However, in the same page, Norton also writes the following: "[w]hat warrants an inductive inference? Answer: facts". Thus, the tension remains.

However, this tension can be resolved by adding some more structure to our understanding of factual warrant. The result is the GFW. According to the GFW, theory and idealisations can (and often do) function as ideal warranting conditions for a given argument, $\mathrm{IWC_R(A)}$. Since ideal warranting conditions are *ideals*, they do not have to be true. Thus, any proposition, independently of its truth value, can play the role of $\mathrm{IWC_R(A)}$. How warranted our rule is in a context depends on how close we are to meeting its ideal warranting conditions. The warrant of the rule – the source of support – is still found in the *facts* about the world.

Take argument [Matter], adapted from (Norton, 2021c, p. 80):

P : In our vicinity of the universe, matter is distributed roughly uniformly in galaxies in a space of constant, possibly zero curvature
C : In all regions of the universe, matter is distributed roughly uniformly in galaxies in a space of constant, possibly zero curvature

[Matter]

This argument can be very naturally understood in terms of EIS:

> EIS([Matter]): if in our vicinity of the universe matter is distributed roughly uniformly in galaxies in a space of constant, possibly zero curvature, then in all regions of the universe matter is distributed roughly uniformly in galaxies in a space of constant, possibly zero curvature
>
> IWC_{EIS} ([Matter]): all regions of the universe are like our vicinity of the universe with respect to the distribution of matter

Notice that IWC_{EIS} ([Matter]) is a reformulation of the cosmological principle – we have obtained it just by filling in the placeholders in IWC_{EIS} in the context of argument [Matter]. Hence, Norton does not need to concede that the cosmological principle (an idealisation) is the warranting fact for argument [Matter]; it constitutes its *ideal* warranting conditions, IWC_{EIS} ([Matter]). And there is nothing wrong with idealisations playing this role. Argument [Matter] is still warranted by the facts about our universe; in particular, the facts about how matter is *actually* distributed in it. The more similar the whole universe is to our vicinity of the universe with respect to the distribution of matter (in the relevant way, given our interests), the more truthlike IWC_{EIS} ([Matter]) will be, the more warranted EIS will be in argument [Matter] (as per the GFW) and the more P will support C in this argument (as per the HTI). Idealisations can play a role in our account of inductive support; as anticipated in the previous section, ideal warranting conditions tell us which features about our world are relevant for the warrant of a given rule. For argument [Matter], it is those features of the world that bear on $v(\text{IWC}_{\text{EIS}}$ ([Matter])). This is an important role for idealisations, but not a *warranting* role.

In this section I have offered a further argument in favour of the GFW. In order to account for the role of idealisations in our understanding of inductive support, Norton allows for idealisations to play a warranting role. This threatens to make his MTI inconsistent, since idealisations are false and, according to Norton, only truths can warrant arguments. However, this tension can be resolved by adding some more structure to our understanding of factual warrant, thus obtaining the GFW. According to the GFW, idealisations can, and often do, play the role of ideal warranting conditions. Since ideal warranting conditions are *ideals*, any statement, independently of its truth value, can play this role. How warranted a rule is in the context of an argument A depends on how close we are to meeting its ideal warranting conditions IWC_{R}(A). Thus, in the spirit of Norton: the warrant is still found in the facts about the world.

2.9 Objections

In this section I address some potential objections to the GFW. Some readers may wonder whether there is a unique schema that can play the role of IWC_R for a given rule R. The answer depends on how we individuate sentence schemas. Strictly speaking, two sentence schemas that differ in any way on the syntax of the template or on the side condition might be considered different sentence schemas. However, this is not an interesting distinction in itself. What is interesting for our purposes is the behaviour of the schemas in terms of the truthlikeness of their instances. Two sentence schemas IWC_R and IWC'_R might differ slightly while $v(IWC_R(A)) = v(IWC'_R(A))$ for all arguments A for which R(A) can be formed. For instance, this is the case for schemas IWC_{EIS} and IWC'_{EIS} given the same side condition:

IWC_{EIS}: all *X*s are like the *X*s in set *S* with respect to $p(X)$

IWC'_{EIS}: with respect to $p(X)$, all *X*s are like the *X*s in set *S*

If one wishes to refer to both of these schemas as different schemas then it is the case that there is more than one schema that can play the role of IWC_R, and this is unproblematic. Instead, I am happy to think about schemas IWC_{EIS} and IWC'_{EIS} as one and the same schema.

However, we must be careful not to take this point too far; there is only so much we can change in a schema before its behaviour in terms of truthlikeness changes. For example, it is not true that any schema that is logically equivalent to a given IWC_R is another possible IWC'_R. We may call two schemas "logically equivalent" if, when their placeholders are filled out under the same circumstances, the instances they generate have the same truth value. However, we are now interested not only in what happens when these expressions are true or false but also in what happens with regards to their truthlikeness. Take the following schema:

U: $p(X)$ is uniform across all *X*s

It is the case that $U(A) \iff IWC_{EIS}(A)$ and, therefore, these schemas are logically equivalent. However, it is not the case that $v(U(A)) = v(IWC_{EIS}(A))$ and, accordingly, it is not the case that $v(EIS(A)) = v(U(A))$. This is easy to see by considering a scenario under which $v(U(A))$ is different from both $v(EIS(A))$ and $v(IWC_{EIS}(A))$. Imagine that there happens to be one allotrope of bismuth, which melts at 5,000ºC, and which represents 1% of the total quantity of bismuth in the universe. Imagine also that all bismuth on earth happens to correspond to this allotrope of bismuth. Take our argument of interest to be argument [Bismuth†]:

P :	The melting temperature for some samples of bismuth is 5,000ºC	[Bismuth†]
C :	The melting temperature of bismuth is 5,000ºC	

Under these circumstances, we obtain the following sentences:

(1) EIS([Bismuth†]): if the observed samples of bismuth melt at 5,000ºC, then all samples of bismuth melt at 5,000ºC.

(2) IWC_{EIS}([Bismuth†]): all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature.

(3) U([Bismuth†]): melting temperature is uniform across all samples of bismuth.

And, correspondingly:

(1) v(EIS([Bismuth†])) is very low, since the melting temperature of the observed samples of bismuth is not representative of the melting temperature of all other samples of bismuth in the universe.

(2) v(IWC$_{EIS}$([Bismuth†])) is very low, since 99% of the samples of bismuth in the universe are not like the observed samples of bismuth with respect to their melting temperature.

(3) v(U([Bismuth†]) is very high, since 99% of all bismuth in the universe melts at the same temperature.

Therefore, since v(EIS([Bismuth†])) $\neq$ v(U([Bismuth†]), U cannot function as IWC'$_{EIS}$, despite it being the case that U(A) $\Longleftrightarrow$ IWC$_{EIS}$(A).

One might still be worried that some schema W that is substantially different from IWC$_R$ could fulfil the condition that v(R(A)) $=$ v(W(A)) for all A for which R(A) can be formed. That is, one might be worried that there is a schema W which is not the result of merely rearranging some terms in IWC$_R$ and such that v(R(A)) $=$ v(W(A)) for all A for which R(A) can be formed. Since there is no general approach to exposing IWC$_R$ for any given rule R, it is hard to argue whether this is possible or not. However, I fail to see why this is problematic; while I am inclined to think that there is no such schema W, if there is then we have identified multiple sets of conditions that we can use as IWC$_R$ and we get to pick our favourite.

A more specific worry might be that one such schema W may be an alternative IWC'$_R$ while assessing v(W(A)) would require us to be able to assess v(C), where C is the conclusion of our argument of interest, argument A. This would obviously be problematic, since we are not able to assess v(C) – this is precisely why we are interested in argument A. However, such a schema would be too informative and therefore it would not be the case that v(R(A)) $=$ v(W(A)).

Let me illustrate this with an example. One could argue that for rule EIS, the following schema may function as IWC'$_{EIS}$:

Z: C

with the side condition that C is the conclusion of our argument of interest.

Suppose now that we are interested in argument [Bismuth†], but that the premise of that argument happens to be false (because of errors in the measurement procedure or for any other reasons). Remember that logic is not concerned with whether our premises are true, but only with what they would tell us about our conclusion if they were true. In particular, assume our universe is such that all chemical elements are perfectly uniform in their physical properties, and that the melting temperature of *all* bismuth is 271ºC. Under these circumstances we obtain the following sentences, where I have included IWC$_{EIS}$([Bismuth†]) for reference:

(1) EIS([Bismuth†]): if the observed samples of bismuth melt at 5,000ºC, then all samples of bismuth melt at 5,000ºC.

(2) IWC$_{EIS}$([Bismuth†]): all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature.

(3) Z([Bismuth†]): the melting temperature of bismuth is 5,000ºC.

And, correspondingly:

(1) v(EIS([Bismuth†])) is very high (maximally high), since in the specified universe all chemical elements are uniform in their physical properties

(2) v(IWC_{EIS}([Bismuth†])) is very high (maximally high), since in the specified universe all chemical elements are uniform in their physical properties

(3) v(Z([Bismuth†]) is very low (maximally low), since in the specified universe no sample of bismuth melts at 5,000ºC

Therefore, since v(EIS([Bismuth†])) $\neq$ v(Z([Bismuth†]), Z cannot function as IWC'_{EIS} despite our initial intuitions. As anticipated, Z contains too much information (information about the specific melting temperature of bismuth) and therefore its truthlikeness is sensitive to features to which it should not be sensitive; schemas Z and EIS are just not connected in the right way, while IWC_{EIS} and EIS are.

2.10 Conclusion

In this chapter I have introduced the GFW, a graded account of factual warrant, and integrated it into the HTI, a hybrid account of inductive support. According to the GFW, given a rule R there is an associated sentence schema IWC_R such that, given any argument A for which R(A) can be formed, v(IWC_R(A)) is a measure of the degree to which rule R is warranted in the context of argument A. Furthermore, according to the HTI, $v(IWC_R(A)) = v(R(A))$.

Both the GFW and the HTI rely on the concept of verisimilitude. However, while verisimilitude is interest-relative, the HTI is interest-independent. I have argued that given a rule R, there is a schema IWC_R such that any measure of v(IWC_R(A)) that we find appropriate given some interests will also be appropriate as a measure of v(R(A)) given the same interests.

In this chapter I have illustrated the GFW and the HTI, by analysing the rule of enumerative induction in terms of the HTI; I have also highlighted two logical virtues of the GFW and the resulting version of the HTI. First, the HTI illuminates the connection between factual warrant and inductive support. For a given set of interests, the facts about the world will determine v(IWC_R (A)). The HTI tells us that $v(IWC_R (A)) = v(R(A))$. R(A) provides information about the relation of inductive support in argument A and, therefore, v(IWC_R (A)) determines the truthlikeness of this information about inductive support. Thus, thanks to the GFW, the HTI provides a clear logical picture of the way in which certain features of the world bear on the degree of support of a given argument.

Second, the HTI can account for the role of idealisations and theory in our understanding of inductive support. While idealisations and theory do play a role, it cannot be a warranting role. This is because, as I argued in chapter 1 following Norton, *facts* warrant arguments, and idealisations and theory are not facts but hypotheses (or bodies of hypothesis), which are either known or suspected to be false. However, the GFW gives enough structure to our understanding of factual warrant to provide a consistent account of the role of idealisations: idealisations (and theory) can, and often do, function as ideal warranting conditions for an argument. Since ideal warranting conditions are *ideals*, any statement, independently of its truth value, can play this role.

Chapter 3

Epistemic Virtues of a Graded Account of Factual Warrant

3.1 Introduction

In chapter 2 I introduced the GFW. I then provided an analysis of the rule of enumerative induction in terms of the HTI, and used that example to highlight two logical virtues of the GFW. In this chapter I provide an analysis of a different rule of induction in terms of the HTI, a rule used for causal inferences in the context of comparative group studies. I then use this example to highlight two epistemic virtues of the GFW.

In this chapter I will talk both of inferences and arguments. My usage of these terms was introduced in chapter 1. An argument is a system of two propositions, a premise P and a conclusion C (both of which may be conjunctions of more simple propositions), where P is supposed to provide support for C. An inference, instead, is composed of beliefs. That is; an inference is a system of two beliefs, a belief in proposition P and a belief in proposition C (both of which may be conjunctions of more simple propositions), where belief in P is supposed to provide support for belief in C. Thus, arguments are for logic to study, and inferences for epistemology. I will now talk about inferences because I will begin to extract some epistemic lessons from the HTI, which is a view on the logic of induction. The relation between logic and epistemology is quite complex, and I will say a bit more about it in the next chapter and in the conclusion of this thesis. For now, we just need to realise that, *ceteris paribus*, stronger arguments will result in stronger inferences. That is: all else being equal, the more P supports C, the more our belief in P will support our belief in C.

Here is how I will proceed. In section 3.2 I articulate my rule of interest as a sentence schema. In section 3.3 I expose the ideal warranting conditions for this rule, and I show that analysing a rule in terms of the HTI can prevent confusions in our understanding of rules of induction. In section 3.4 I argue that a rule's ideal warranting conditions can function as a methodological guide for researchers to ensure strong inferences. This is the first epistemic virtue of the GFW that I discuss. In section 3.5 I argue that a rule's ideal warranting conditions also function as a conceptual guide to the factors that are relevant in assessing an inductive inference's strength. This is the second epistemic virtue of the GFW that I discuss. In section 3.6 I address some objections. Section 3.7

concludes.

3.2 Making Our Rule Explicit

This section, together with the next one, constitute the analysis of a rule of induction from the perspective of the HTI. Analysing a rule of induction from the perspective of the HTI means expressing it as a sentence schema R, and exposing the corresponding schema IWC_{R}. My case study of choice is one type of causal inference made in the context of comparative group studies. Comparative group studies are studies in which the average value of some feature of subjects is compared among different groups of subjects. In interventional comparative group studies researchers have control over the assignment of subjects into the groups under comparison, which is not the case in observational comparative group studies. In this section I make explicit a particular rule of induction I am concerned with in the context of causal inference in comparative group studies. I will call this rule CIS because it captures a particular understanding of Causal Inference under a Strong interpretation. In the next section I will expose the ideal warranting conditions for this rule, IWC_{CIS}.

Let's first make explicit the type of inferences I am concerned with. Randomized Controlled Trials (RCTs) are set up in order to determine the causal relation between two features of a study population: an intervention or treatment, represented by the independent variable T, and an outcome, represented by the dependent variable Y. Any factor other than the intervention that has a causal impact on the outcome is called a "covariate" (Deaton & Cartwright, 2018) or a "confounder" (Fuller, 2019). For simplicity, let us assume that the intervention of interest is dichotomous, as is often the case, so that either $T = 0$ (the intervention is absent) or $T = 1$ (the intervention is present). In order to set up an RCT, researchers randomly assign the individuals in the study population to one of two groups: the control group (for whom $T = 0$) or the treatment group (for whom $T = 1$). Researchers then collect evidence about Y for every individual, and try to determine the causal impact of T on Y.

I will rely on the following example so we can explore causal inferences from RCTs in more detail:

> Imagine a researcher interested in the impact that some job training programs (T) have on labor market outcomes (Y), typically wages or employment-status. Labor market outcomes are related to many variables, such as education (schooling), work experience, gender, and family networks. Let us formalize these relations with a simple linear model: $Y_i = \beta T_i + \boldsymbol{X}_i' \boldsymbol{\gamma}$. Subscript i refers to the individual, T to attending (or not) a job training program, and β to the impact of T on Y. Vectors $\boldsymbol{X}$ and $\boldsymbol{\gamma}$ refer, respectively, to all the minimally sufficient factors (apart from T) that affect labor market outcomes and their respective impact on them.[1] (Larroulet Philippi, 2022, p. 153)

[1] Several assumptions are being made in this example, following Larroulet Philippi. I assume that the impact of T on Y is the same for all individuals (so $\beta_i = \beta$ for all i). I am ignoring Fuller's (2019) distinction between "directly causal" and "association-causal" confounders. I am also assuming that causes act deterministically. These assumptions allow us to set up a simple example, but they are not relevant for the points that I will make. The upcoming analysis can be repeated for any different set of assumptions.

I will call this example "study S_1". For current purposes, I will take the labor market outcome of interest, Y, to be the annual income. Notice that, since vector $\boldsymbol{X}$ contains, by stipulation, all the minimally sufficient factors that affect Y, apart from T, then β represents the causal impact of T on Y.

It is common to understand causal inferences from RCTs as following the rule that "correlation is a sign of causation" or, in its strongest version, that "correlation implies causation". This rule captures the inferences that are made from observations of correlation between T and Y to the claim the T is responsible for the effects on Y. We can see references to this rule in classic accounts of causal inference from RCTs, like that of Papineau (1994) or Cartwright (2007). Cartwright's ideal RCTs (Cartwright, 2007), in particular, are understood as following the strong rule that "correlation implies causation". This is a good but vague characterization of the different types of inferences that researchers make from the evidence obtained from RCTs. It will be an interesting exercise to expose in more detail one of the rules of induction that is often followed in the context of RCTs, and its ideal warranting conditions.

Let's make our rule of interest explicit then. In some common inferences from RCTs, the researchers want to attribute the total effect on Y to the intervention T. This is useful because it allows the researchers to quantify the causal impact of T on Y. In such scenarios, the researchers are applying the rule that "all observed effects on Y are due to T", which can also be expressed in the form of the following schema:

CIS: $\beta_S = E(Y_i|T = 1)_S - E(Y_i|T = 0)_S$

The only placeholder in this schema is S, which must be filled out with the identifier of the comparative group study that our argument is about. The rest of symbols are significant terms, which must be interpreted as follows: Y refers to the outcome of study S, T refers to the treatment of study S, β_S refers to the causal impact of T on Y in study S, $E(Y_i|T = 1)_S$ refers to the expected value of Y_i for the participants in the treatment group of study S, and $E(Y_i|T = 0)_S$ refers to the expected value of Y_i for the participants in the control group of study S. This is the side condition of schema CIS. Notice that the only placeholder in rule CIS is filled in with information about our argument of interest A and, therefore, I will refer to instances of CIS as CIS(A).

CIS allows researchers to make inferences of the following type:

$$\begin{array}{ll} P: & E(Y_i|T = 1)_S - E(Y_i|T = 0)_S = k \\ C: & \beta_S = k \end{array} \tag{3.1}$$

For instance, in the case of study S_1, researchers may have evidence that $E(Y_i|T = 1)_{S_1} - E(Y_i|T = 0)_{S_1} = 10,000€$, and may be interested in making the following inference:

$$\begin{array}{ll} P: & E(Y_i|T = 1)_{S_1} - E(Y_i|T = 0)_{S_1} = 10,000€ \\ C: & \beta_{S_1} = 10,000€ \end{array} \tag{Job}$$

Inference [Job] is, as it stands, an inductive inference. The evidence (P) for these researchers is the difference in expected values of Y_i (annual income) between both arms of trial S_1. The conclusion (C) for these researchers is that this difference is the value of β_{S_1}. The conclusion contains information not contained in the premises, and thus the inference is inductive.

CIS([Job]) describes the relation of support in argument [Job]. CIS([Job]) is the statement $\beta_{S_1} = E(Y_i|T = 1)_{S_1} - E(Y_i|T = 0)_{S_1}$ which, when interpreting the symbols according to the side condition, is equivalent to the following statement: "the causal impact of the treatment of study S_1 on the outcome of study S_1, in study S_1, is equal to the the expected value of the outcome of study S_1 for the participants in the treatment group of study S_1 minus the expected value of the outcome of study S_1 for the participants in the control group of study S_1". Notice that this statement contains definite descriptions like "the treatment of study S_1" or "the outcome of study S_1", which are proper. If desired, these can be substituted for coreferential terms like "the job training program" or "the annual income", correspondingly. In that case, we obtain the following statement: "the causal impact of the job training program on the annual income, in study S_1, is equal to the the expected value of the annual income for the participants in the treatment group of study S_1 minus the expected value of the annual income for the participants in the control group of study S_1". CIS([Job]) describes the relation of support in argument [Job] because it tells us that P implies C in argument [Job].

This sounds deductive. Indeed, whenever CIS([Job]) is fully warranted, we can turn [Job] into a deductive argument by adding CIS([Job]) as an explicit premise. This is a limiting case – what Cartwright calls "ideal RCTs" (Cartwright, 2007) – and is no challenge to the view I am providing here.[2] As we will see, in real RCTs, CIS(A) will (in general) not be fully warranted, so it will be, strictly speaking, false. Hence, we will not add CIS(A) as a premise, and our arguments will remain inductive. As I will show throughout this chapter, the HTI provides the right framework to understand what is happening in these arguments.

In this section I have made rule CIS explicit. Rule CIS is the schema $\beta_S = E(Y_i|T = 1)_S - E(Y_i|T = 0)_S$, with the corresponding side condition. This rule tells us that P implies C in arguments like [Job]. However, it is widely acknowledged that this is not always true, and the strength of such arguments is variable. In the next section I expose IWC_{CIS} and show how it allows us to capture and understand the strength of such arguments.

3.3 Exposing the Ideal Warranting Conditions of Our Rule

After making rule CIS explicit, we now have to determine IWC_{CIS}. That is, we have to find a schema IWC_{CIS} such that, given any argument A for which CIS(A) can be formed, $v(\text{IWC}_{\text{CIS}}(\text{A})) = v(\text{CIS}(\text{A}))$. That is the goal of this section.

Developing the right-hand side of rule CIS will be useful:

$$\begin{aligned} E(Y_i|T = 1)_S - E(Y_i|T = 0)_S &= E(\beta T_i + \boldsymbol{X}_i'\gamma|T = 1)_S - E(\beta T_i + \boldsymbol{X}_i'\gamma|T = 0)_S \\ &= \beta_S + E(\boldsymbol{X}_i'\gamma|T = 1)_S - E(\boldsymbol{X}_i'\gamma|T = 0)_S \end{aligned} \tag{3.2}$$

Isolating β_S from expression 3.2, we obtain

$$\beta_S = [E(Y_i|T = 1)_S - E(Y_i|T = 0)_S] - \left[E(\boldsymbol{X}_i'\gamma|T = 1)_S - E(\boldsymbol{X}_i'\gamma|T = 0)_S\right] \tag{3.3}$$

[2] I address this concern in more depth in §3.6.

We can now see that rule CIS follows from the general expression 3.3 when

$$E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S = 0 \tag{3.4}$$

or, equivalently, when

$$E(\boldsymbol{X}_i'\gamma|T=1)_S = E(\boldsymbol{X}_i'\gamma|T=0)_S \qquad \text{NBAL}$$

I call this schema NBAL because it is a particular type of balance assumption between both arms of a trial, what I call the "New BALance assumption", in contrast with older versions of this condition. I will discuss the contrast between the old and new versions of the balance assumption below. Let me now specify the side condition for schema NBAL. The only placeholder in this schema is S, which must be filled out with the identifier of the comparative group study that our argument is about. The rest of the symbols are significant terms, which must be interpreted as follows: T refers to the treatment of study S, vector $\boldsymbol{X}$ refers to all the minimally sufficient factors (apart from T) that affect the outcome of study S, vector γ refers to the impact of these factors on the outcome of study S and, therefore, $E(\boldsymbol{X}_i'\gamma|T=1)_S$ refers to the expected value of $\boldsymbol{X}_i'\gamma$ for the participants in the treatment group of study S, and $E(\boldsymbol{X}_i'\gamma|T=0)_S$ refers to the expected value of $\boldsymbol{X}_i'\gamma$ for the participants in the control group of study S. In short, schema NBAL tells us that the sum of all confounders' effects is balanced between both arms of the trial.

Schema NBAL is IWC$_{\text{CIS}}$. In order to show this, we need to show that $v(\text{CIS(A)}) = v(\text{NBAL(A)})$ for any argument A such that CIS(A) can be formed. From expression 3.3 we can obtain the following expression, which will be useful for our purposes;

$$\beta_S - [E(Y_i|T=1)_S - E(Y_i|T=0)_S] = -\left[E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right] \tag{3.5}$$

Next, let's show that $v(\text{CIS(A)}) = v(\text{NBAL(A)})$ independently of our interests.

3.3.1 First set of interests

Let's start by determining a measure of v(NBAL(A)). NBAL(A) is an equality. For instance, NBAL([Job]) is the equality $E(\boldsymbol{X}_i'\gamma|T=1)_{S_1} = E(\boldsymbol{X}_i'\gamma|T=0)_{S_1}$. Thus, we need to determine a measure of truthlikeness of an equality. Suppose that the difference between both sides of the equality is representative of the magnitude of our concern about this difference. Then, an obvious candidate measure for the truthlikeness of this equality is the absolute value of the difference between both sides of the equality, with a negative sign to maintain the convention that larger values of v indicate more truthlikeness, i.e.,

$$-|L-R| \tag{3.6}$$

where L refers to the left-hand side of the equality and R to the right-hand side.

Thus,

$$v(\text{NBAL(A)}) = -\left|E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right| \tag{3.7}$$

As desired, if expression 3.6 is a good measure of v(NBAL(A)) it will also be a good measure of v(CIS(A)). This is easy to check. CIS(A) is also an equality. For instance, CIS([Job]) is the equality $\beta_{S_1} = E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1}$. Since our interests are the same as before, we still care about the difference between both sides of an equality proportionally to the size of those differences. Thus, an appropriate measure for v(CIS(A)) would be $-|L-R|$ and, therefore,

$$v(\text{CIS(A)}) = -\left|\beta_S - [E(Y_i|T=1)_S - E(Y_i|T=0)_S]\right| \tag{3.8}$$

Now, if we take expression 3.5 and take the absolute value of each side, we obtain:

$$\left|\beta_S - [E(Y_i|T=1)_S - E(Y_i|T=0)_S]\right| = \left|E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right| \tag{3.9}$$

and, therefore, from expressions 3.7, 3.8 and 3.9, we see that, for measure 3.6, $v(\text{CIS(A)}) = v(\text{NBAL(A)})$.

3.3.2 Second set of interests

Even if we change our interests, $v(\text{CIS(A)}) = v(\text{NBAL(A)})$. Say that, now, we care about bigger differences between both sides of an equality much more than we care about smaller differences, so that the following is an adequate measure of its truthlikeness:

$$-(L-R)^2 \tag{3.10}$$

Thus,

$$v(\text{NBAL(A)}) = -\left[E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right]^2 \tag{3.11}$$

and

$$v(\text{CIS(A)}) = -\left[\beta_S - [E(Y_i|T=1)_S - E(Y_i|T=0)_S]\right]^2 \tag{3.12}$$

Now, if we take expression 3.5 and square each side, we obtain:

$$\left[\beta_S - [E(Y_i|T=1)_S - E(Y_i|T=0)_S]\right]^2 = \left[E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right]^2 \tag{3.13}$$

and, therefore, from expressions 3.11, 3.12 and 3.13, we see that, for measure 3.10, it is also the case that $v(\text{CIS(A)}) = v(\text{NBAL(A)})$.

3.3.3 Inadequate measures of truthlikeness

While there is some freedom in choosing a measure of the truthlikeness of an expression, not anything goes. Here is an example; given an equality, $a = b$, suppose that we want to use the following as a measure of the truthlikeness of this expression:

$$-\frac{|L-R|}{L} \tag{3.14}$$

The minus sign has been added to respect the convention that larger values of our measure of verisimilitude represent more closeness to the truth. So, according to measure 3.14,

$$v(a=b) = -\frac{|a-b|}{a} \tag{3.15}$$

Measure 3.14 might seem appropriate in order to capture the following interests; the difference between both sides of an equality is relevant proportionally to the size of (one of) those terms.

However, measure 3.14 is an inadequate measure of truthlikeness of an equality, since it would yield conflicting results when applied to equivalent expressions. Expression $a = b$ is equivalent to expression $b = a$. However, if $a \neq b$, measure 3.14 would yield different results when applied to both expressions, which is problematic, since expression $a = b$ is as truthlike as expression $b = a$.

It seems that this could be solved by suggesting a slightly different measure, where the absolute value of the difference between both sides of an equality is divided by the absolute value of the average of both sides:

$$-\frac{|L-R|}{\frac{|L+R|}{2}} \tag{3.16}$$

However, while measure 3.16 does yield the same result when applied to expressions $a = b$ and $b = a$, it is still inadequate. This measure does not yield the same result when applied to the equivalent expression $a - b = 0$. Furthermore, measure 3.16 does not behave well when one of the sides of our equality of interest is 0, in which case it returns the constant value -2. However, expression $10^{-30} = 0$ seems to be more truthlike than expression $10^{30} = 0$, but measure 3.16 cannot capture this, since it would assign the same value to the truthlikeness of both expressions.

In contrast, measures 3.8 and 3.12, introduced above, behave adequately. For instance, take the following expressions:

$$E(\boldsymbol{X}_i'\gamma|T=1)_{S_1} = E(\boldsymbol{X}_i'\gamma|T=0)_{S_1} \tag{3.17}$$

$$E(\boldsymbol{X}_i'\gamma|T=0)_{S_1} = E(\boldsymbol{X}_i'\gamma|T=1)_{S_1} \tag{3.18}$$

$$E(\boldsymbol{X}_i'\gamma|T=1)_{S_1} - E(\boldsymbol{X}_i'\gamma|T=0)_{S_1} = 0 \tag{3.19}$$

Expressions 3.17, 3.18 and 3.19 are all equivalent, and correspond to $\mathrm{IWC}_{\mathrm{CIS}}([\mathrm{Job}])$. Measures 3.14 and 3.16 would yield different results when applied to these expressions, but applying measure 3.6 to expressions 3.17, 3.18 and 3.19 yields the same result, and so does measure 3.10, as desired.

Asking a measure of truthlikeness to deliver the same result when applied to equivalent expressions is just one of the many adequacy conditions we might impose on such measures. Determining what constitutes an *adequate* measure of truthlikeness is outside the scope of this thesis (see (Niiniluoto, 1998) for a survey of some proposed measures of truthlikeness and adequacy conditions on such measures).

3.3.4 Bringing clarity

That NBAL is $\mathrm{IWC}_{\mathrm{CIS}}$ is quite an intuitive result. Rule CIS tells us that all observed effects on Y are due to T. $\mathrm{IWC}_{\mathrm{CIS}}$ tells us that the truthlikeness of this depends on how balanced the sum of all confounders' effects on Y is between both arms of the trial. However, despite being an intuitive result, we should not underestimate the need for precisely expressing our rule of interest and its ideal warranting conditions. These expressions need not be formal, as they are here, but they must be precise and they must be made explicit. Doing this work allows for better discussion and prevents misunderstandings. I now provide two examples of the kind of misunderstandings this analysis prevents.

First, making explicit the rule we are interested in is useful, for instance, in order to prevent confusions between similar but different rules. The vague rule "correlation implies causation" can be interpreted in different ways. I have focused here in a common strong interpretation of this rule, which is followed when researchers wish to attribute all observed effects on Y to intervention T. In these cases, our rule is CIS. However, in other cases, researchers may only be interested in establishing that T has *some* causal impact on Y (Fuller, 2019, 919).[3] In those situations, the rule being applied is a weaker one, viz, the schema

$$\beta_S > 0 \text{ iff } E(Y_i|T=1)_S - E(Y_i|T=0)_S > 0 \tag{3.20}$$

where the only placeholder is S, and the rest of the terms are to be understood as previously specified.

Notice that this weaker rule is also a version of the claim that "correlation implies causation". This weaker rule, of course, requires weaker ideal warranting conditions. In particular, the ideal warranting conditions for this rule claim that the imbalance in the overall confounders' effect on the outcome is smaller than the total observed difference in outcome, which can be expressed as the schema

$$E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S < E(Y_i|T=1)_S - E(Y_i|T=0)_S \tag{3.21}$$

where the only placeholder is S, and the rest of the terms are to be understood as previously specified.

[3] I ignore here cases of prevention, where the presence of T prevents the outcome Y. A different rule, with different ideal warranting conditions, can be determined for this type of inference.

Second, being explicit about which rule of induction we wish to use allows us to do the work required to expose its ideal warranting conditions. Most classic literature (e.g. Byar et al., 1976; Cartwright, 1979; Papineau, 1994; Schwartz et al., 1980) takes the condition that *every* confounder be balanced among study groups to be a necessary requirement for causal inference in comparative group studies, and this is still an influential view today.[4] Worrall's criticism of this view (Worrall, 2002) has, correspondingly, also become influential. I will call this condition the "old balance assumption".

However, some authors have argued that the old balance assumption is too strict an ideal (Fuller, 2019; Larroulet Philippi, 2022), and it is not necessary for causal inference. This is correct, and the analysis just performed in this section allows us to clearly see why. Rule CIS, the strongest interpretation available of the claim "correlation implies causation", requires weaker ideal warranting conditions (IWC_{CIS}) than the old balance assumption. IWC_{CIS} is NBAL, the "new balance assumption": it is not balance in every confounder that is required, but only balance in the sum of the confounder's effects. In the next section I discuss why this is a relevant difference. We can see that the new balance assumption is a weaker requirement than the old balance assumption because, while it is possible to satisfy the new balance assumption without satisfying the old balance assumption, it is not possible to satisfy the old balance assumption without satisfying the new balance assumption. In sum, the confusion in classic literature is a confusion about the ideal warranting conditions that our rule requires, and this can be clarified by analysing our rules of induction from the perspective of the HTI.

In the previous section I made rule CIS explicit, and in this section I have exposed the ideal warranting conditions for this rule. While I hope that this exercise is illuminating in itself, I have argued that this type of analysis is useful too. On the one hand, analysing our rules of induction in terms of the HTI forces us to express them precisely, in the form of a sentence schema, which prevents confusions. Different rules of causal inference are captured by the same slogan that "correlation implies causation", but once we express our rule of interest in the form of a sentence schema, the differences between those rules become explicit and clear. On the other hand, the HTI also forces us to expose the ideal warranting conditions for our rule of interest, and it tells us how we should go about it. Given rule CIS, our schema IWC_{CIS} must be such that $v(\mathrm{CIS}(A)) = v(\mathrm{IWC_{CIS}}(A))$. In this way we realise that the old balance assumption, requiring balance in every confounder to allow for applications of rule CIS, is too strong an ideal; IWC_{CIS} is NBAL, which only requires balance in the total contribution of the confounders. In the next section I develop a further lesson, i.e., that exposing a rule's ideal warranting conditions can act as a methodological guide too.

3.4 Getting Close to the Ideal Warranting Conditions

In this section I argue that a rule's ideal warranting conditions can function as a methodological guide for researchers to ensure strong inferences. The debate around the importance of randomization in comparative group studies is a case in point. Urbach (1985) and Worrall (2002, 2007) have famously argued that "there is no cause to randomize" (to use Worrall's words) in the context

[4] Each of these authors makes this point in different terms. Cartwright (1979), for instance, makes her point in terms of "state descriptions", that is, any possible arrangement of all alternative causal factors (besides our factor of interest).

of subject allocation for comparative group studies. Larroulet Philippi (2022) has responded that there is. I agree with Larroulet Philippi, and I will argue here that the HTI provides a clear logical picture of why this is so.

Let's now fix our rule of interest and its ideal warranting conditions. Recall that comparative group studies are studies in which the average value of some dependent variable is compared among different groups of subjects, typically a control group and a treatment group. I will focus for now on interventional (as opposed to observational) comparative group studies, in which the researchers have control over the subjects' assignment to the experimental groups. These studies are set up in order to determine the causal impact of some intervention. Under these circumstances we will most often want to use rule CIS, which allows us to obtain a value for β_S (as opposed to rule 3.20), so that our ideal warranting conditions are IWC_{CIS} (that is, NBAL).

3.4.1 To randomize or not to randomize

The key point of this section is the following: knowing that $v(\text{CIS(A)}) = v(\text{IWC}_{\text{CIS}}(\text{A}))$ for a given argument A is a methodological guide. CIS(A) tells us that *P* implies *C* in arguments like [Job], and this is more truthlike the greater $v(\text{IWC}_{\text{CIS}}(\text{A}))$. The more truthlike it is that *P* implies *C*, the more *P* supports *C* and, therefore, the stronger our inference from *P* to *C*.[5] Thus, our inferences will be stronger the greater $v(\text{IWC}_{\text{CIS}}(\text{A}))$ is, i.e., the more balanced the overall contribution of the confounders is among study groups. Therefore, in order to maximize the strength of our inferences we shall get as close to $v(\text{IWC}_{\text{CIS}}(\text{A}))$ as possible.[6] A question we can ask as researchers, then, is: what can we do to get close to $\text{IWC}_{\text{CIS}}(\text{A})$?

One might say that we should match the study groups for all known observable confounders.[7] To go back to our example; if we know that market outcomes, like annual income (Y), are not only affected by our job training program of interest (T) but also by work experience and gender, we should construct our study groups by making sure that the work experience and gender of participants is balanced among them. Every known observable confounder that is controlled for in this way takes us one step closer to $\text{IWC}_{\text{CIS}}(\text{A})$.

Even if we agree that matching is beneficial, the following question remains: after controlling for all known observable confounders, is there any extra value in *randomly* assigning the subjects to these groups?[8] Many practitioners think so, and this is why they carry out RCTs. However, Urbach (1985) and Worrall (2002) disagree. The following fragment from Worrall is representative of this criticism:

> Even if there is only a small probability that an individual factor is imbalanced, given that there are indefinitely many possible confounding factors, then it would seem to follow that the probability that there is some factor on which the two groups are

[5] To say that *P* implies *C* is to say that the degree of support is maximal; the claim that "the degree of support is maximal" will be more truthlike the higher the degree of support. Recall also that, all else being equal, the more *P* supports *C*, the more our belief in *P* supports our belief in *C*.

[6] Other factors, like ethical or economic considerations, also influence our methodological choices. I here ignore these factors and focus only on the impact that our methodological choices have on the strength of our inferences.

[7] Although some argue that this is an unnecessary complication on randomization (e.g. Peto et al., 1976).

[8] This is the central question that Larroulet Philippi (2022) tackles.

> imbalanced (when remember randomly constructed) might for all anyone knows be high. (Worrall, 2002, S324)

According to these critics, there is no extra value in randomizing because randomization does not take us closer to the *old* balance assumption, since the probability that some confounder is imbalanced remains high even after randomizing. But we shouldn't care about imbalances in individual confounders, but about imbalance in the sum of all confounders' effects on our outcome. Now that we have phrased the debate in terms of the HTI, the answer to this criticism is straightforward: it doesn't matter whether randomization takes us closer to the *old* balance assumption, but whether it takes us closer to the *new* balance assumption NBAL(A), which is $\text{IWC}_{\text{CIS}}(\text{A})$.

And randomization does take us closer to the new balance assumption (Fuller, 2019; Larroulet Philippi, 2022). The new balance assumption requires that the sum of the confounders' effects be the same in both study groups. While randomizing does not ensure that any single confounder is balanced among study groups (as Worrall points out), it does ensure that the imbalanced contributions of individual confounders are randomly distributed among study groups. This random distribution of individual imbalances is better than our alternative, which is to assign individuals to study groups by using an unsuspicious variable U, which "is a variable, not randomly generated, about which the researcher has no positive evidence that it will produce baseline imbalance" (Larroulet Philippi, 2022, 159).[9] In the best case scenario, where a chosen U is actually unrelated to any confounder, then the assignment method is as good as randomization. However, there are two ways in which picking a U may go wrong (Larroulet Philippi, 2022, 160–63): a researcher may be ignorant that a chosen U is actually systematically related to some confounders, or the researcher might be unaware that her reasons for picking a particular U are related to some confounders. These are problems that do not arise with randomization. Therefore, there is cause to randomize because it is the subject allocation method that ensures we are as close to $\text{IWC}_{\text{CIS}}(\text{A})$ as possible, i.e., it maximizes $v\,(\text{IWC}_{\text{CIS}}(\text{A}))$.

The point I've just made is important in two ways. First, it frames and emphasizes a key discussion we must have: we must make an effort to explore and explicitly state the ideal warranting conditions of the rules of induction we wish to use. This discussion, in the context of causal inference, has been tackled by several authors (e.g. Britton et al., 1999; Cartwright, 2007; Fuller, 2019; Larroulet Philippi, 2022; Papineau, 1994; Yusuf, Held, Teo, & Toretsky, 1990), and we can understand their diverse arguments from the common framework that the HTI provides.

Second, we must notice that this debate is not only conceptually interesting but has practical implications too. As discussed above, the decision whether to randomize or not depends (at least in part) on whether randomization takes us closer to the ideal warranting conditions of the rule we wish to apply, and, in order to determine that, we must first expose the right ideal warranting conditions. There are other reasons in favour of randomizing (Larroulet Philippi, 2022, §7), but the fact that it brings us closer to $\text{IWC}_{\text{CIS}}(\text{A})$ is a crucial one. In general, while we may have other reasons to make some methodological choices in a given study, we can't deny the importance (and maybe the priority) of maximizing the strength of our inferences. Therefore, the ideal warranting conditions of our rule of choice can be an important methodological guide.

[9] Our alternative is *not* to keep matching for confounders; remember, I am assuming that we have already matched for all known observable confounders.

3.4.2 Other methodological choices

Randomization is not the only tool we can use to get close to the new balance assumption. I have already mentioned above that matching for known observable confounders before randomizing, for example, is also important. This is a helpful strategy because it brings us closer to the new balance assumption. Since known observable confounders are matched on an individual basis, this strategy would also bring us closer to the old balance assumption, but this is irrelevant and should not cause confusion. We do not match known confounders to get close to the old balance assumption, but to the new one.

Blinding is also important. Since subject allocation happens before the treatment is implemented, randomization can only help us avoid imbalance in the sum of confounders' effects at *baseline* (that is, before the intervention). However, further confounders may sweep in after the allocation of subjects has taken place. A patient's self-assessment of her health condition (which might be the output of a comparative group study), for instance, can be influenced by her knowledge that she is in the treatment group as opposed to the control group. Blinding refers to the concealment of subject allocation from one or more groups of individuals participating in a comparative group study (participants, clinicians, data collectors, outcome adjudicators and data analysts). Blinding of the different groups involved in an RCT helps reduce various biases (Karanicolas, Farrokhyar, & Bhandari, 2010, p. 346) of which performance bias and detection bias might be the most notable. Biases are a specific kind of confounder, since they are causes of a study outcome other than the intervention. By blinding we minimize the effect of some biases, which are confounders, so we get closer to the new balance assumption.

3.4.3 Concluding remarks

In conclusion, a rule's ideal warranting conditions can function as a methodological guide to ensure strong inferences. I have argued that some popular techniques in comparative group studies, like matching, randomization or blinding, are appropriate precisely because they bring us closer to the new balance assumption, which describes the ideal warranting conditions of the rule of induction we wish to apply in that context.[10] This perspective can be used to provide a sharper defence against some critics of a particular technique, like Urbach or Worrall.

However, a rule's ideal warranting conditions can only function as a methodological guide when there are steps we can take to get closer to them. This is not always the case. Recall argument [Bismuth], for which IWC_{EIS} ([Bismuth]) is the statement that "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature". There is nothing we can do to get closer to these ideal warranting conditions – there are no steps we can take, for instance, to increase the uniformity in the melting temperature of bismuth. In this case, then, our rule's ideal warranting conditions will not function as a methodological guide. Nonetheless, ideal warranting conditions *always* function as a conceptual guide to the factors that are relevant in assessing an inductive inference's strength. I develop this point in the next section.

[10] Fuller (2019, p. 920) neatly expresses this same point, in his own terms, for the case of comparative group studies.

3.5 Judging How Close we Got

In the previous section I have argued that a rule's ideal warranting conditions can function as a *methodological* guide to ensure strong inferences. In this section I argue that, furthermore, a rule's ideal warranting conditions also function as a *conceptual* guide to the factors that are relevant in assessing an inductive inference's strength. This is the epistemic version of the logical point I made in §2.7, and I hope it provides further clarification. Thus, exposing a rule's ideal warranting conditions allows us to (1) make the methodological choices available (if any) to get as close to them as possible, and (2) judge how close we got so we can properly assess the strength of our inferences.

Take rule CIS and its ideal warranting conditions, IWC_{CIS}, for instance. CIS([Job]) tells us that *P implies C* in argument [Job] – it tells us that the degree of support in that argument is maximal. This is only true, however, when rule CIS is fully warranted, i.e., when IWC_{CIS}([Job]) is true. When IWC_{CIS}([Job]) is false, $v(IWC_{CIS}([Job]))$ will tell us how accurate rule CIS is in the context of argument [Job] and, therefore, how accurate it is that *P* implies *C* in that argument; the more accurate this is, the stronger the argument and, therefore, the stronger the inference.

IWC_{CIS}, then, is a conceptual guide to the factors which are relevant in assessing our inferences' strength. IWC_{CIS} is a conceptual guide to such factors in the following sense: it tells us that the factors relevant in assessing the strength of a given inference A are those which are relevant for the truthlikeness of IWC_{CIS}(A). The size of the study groups or the allocation method of participants are relevant factors in assessing the strength of inference [Job] because they bear on the truthlikeness of IWC_{CIS}([Job]), since *large* study groups with *randomly* allocated subjects bring us closer to it.

The role of IWC_{CIS} as a conceptual guide becomes even more clear when analysing a comparative group study that is already done. No methodological choices can be made now, so IWC_{CIS} cannot play the role described in the previous section. However, we must still critically assess the level of confidence we should place in the conclusions of the study. Say we are interested in a particular study in which the researchers are trying to quantify the causal impact of a given variable over the study outcome (the value of β_S). We can understand the argument of interest A, relating the evidence and the hypothesis of those researchers, from the perspective of rule CIS, so the corresponding ideal warranting conditions are IWC_{CIS}(A). The more truthlike IWC_{CIS}(A), the stronger the argument and the more confident we should be in the conclusion of the study.

These are, indeed, the kinds of considerations we make when assessing the strength of causal inferences in comparative group studies. We value large study groups or randomly allocated subjects because these strategies bring us closer to the new balance assumption. Similarly, both in interventional and observational studies, we use background knowledge to ensure that known observable confounders are matched. Furthermore, in observational studies, we use background knowledge to check whether there are known unobservable confounders, since there is no randomization to balance their overall contribution here. We then place more confidence in those studies in which a good background knowledge suggests no unobservable confounders (Larroulet Philippi, 2022, 164–65; Fuller, 2019, 923). We value all these factors, and place more confidence in the conclusions of studies in which they are present, because they bring us closer to the new balance assumption. The HTI gives us a clear logical picture of why these considerations are right, and, correspondingly, it provides a clear defence against critics of such considerations.

How to weigh all these relevant factors is a more complicated matter. Is randomization more or less important than having a quite complete background knowledge that suggests no unobservable confounders? This is a hard question to answer.[11] My contribution here is not to provide a method to bundle all these considerations together into an ultimate quality indicator. My contribution is to provide an underlying logical framework that brings to light why certain factors are relevant in assessing the strength of an inference, and allows us to work out what these factors are. For arguments following rule CIS, the relevant factors in assessing the strength of the corresponding inferences are those which are relevant for the truthlikeness of $\text{IWC}_{\text{CIS}}(\text{A})$.

This lesson is general, and it applies equally to all rules and their ideal warranting conditions. I already anticipated this point in §2.7. Recall inference [Salt]:

P :	This sample of radium chloride is monoclinic	[Salt]
C :	All samples of radium chloride are monoclinic	

Norton tells us that what matters for the strength of this argument is whether or not polymorphism interferes. However, the MTI cannot provide an account of why this is what matters in this case. The HTI, instead, provides a general response. IWC_{EIS} ([Salt]) functions as a conceptual guide to the factors that are relevant in assessing the strength of argument [Salt]. IWC_{EIS} ([Salt]) is the claim that "all samples of radium chloride are like the observed sample of radium chloride with respect to their crystallographic form". Polymorphism is defined as the condition in which a solid chemical compound can exist in more than one crystallographic form. Hence, by definition, polymorphism captures exactly the feature of radium chloride that can make IWC_{EIS} ([Salt]) less truthlike. This is why polymorphism is what matters in assessing the strength of argument [Salt], and this is how IWC_{EIS} ([Salt]) functions as a conceptual guide to the factors that are relevant in assessing the strength of our corresponding inference.

In this section I have argued that a rule's ideal warranting conditions always function as a conceptual guide to the factors which are relevant in assessing the strength of our inferences. The HTI provides a logical framework to understand which factors should impact our confidence in the conclusion of a given inference, and why this is so. In the case of causal inference in comparative group studies, factors such as the presence or absence of blinding, or the size of the study groups, are relevant in assessing the strength of our inferences because they bear on the truthlikeness of $\text{IWC}_{\text{CIS}}(\text{A})$.

3.6 Objections

In this section I briefly address some potential objections to the analysis I just introduced. First, some readers might wonder why I'm treating my case study from the perspective of inductive logic at all. Since there are plenty of numerical methods already available to researchers in the context of causal inference, it might seem that any issue there can be solved with the right mathematical tools. However, no matter how far numerical methods can take us, I am still interested in the

[11] Or perhaps impossible (see Stegenga, 2015).

relation between the researchers' evidence and their conclusions. That relation is one of inductive support, since the researchers' evidence does not imply their conclusion, so this is a relation for inductive logic (and epistemology) to capture.

One could still argue that, given the mathematical relations between variables that I have provided in §3.2, the arguments I am concerned with are, if anything, deductive. But they are not. Recall argument [Job]:

$$\begin{array}{rl} P: & E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} = 10,000€ \\ C: & \beta_{S_1} = 10,000€ \end{array} \qquad \text{[Job]}$$

As it stands, this is an inductive argument, since P does not entail C. But we could, for example, add CIS([Job]) as an explicit premise (P_1) to make it a deductive argument:

$$\begin{array}{rl} P_1: & \beta_{S_1} = E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} \\ P_2: & E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} = 10,000€ \\ C: & \beta_{S_1} = 10,000€ \end{array} \qquad (3.22)$$

However, this deductive argument will not be sound in real scenarios because P_1 will be false, since the total contribution of all confounders will not be balanced.

This may raise a different concern. For real (as opposed to ideal) RCTs, deductive logic actually tells us that C in argument [Job] is false. Let's see why. In real scenarios, the total contribution of all confounders will not be balanced, i.e.,

$$E(\boldsymbol{X}_i'\gamma|T=1)_{S_1} - E(\boldsymbol{X}_i'\gamma|T=0)_{S_1} \neq 0 \qquad (3.23)$$

Now recall the general expression for β, expression 3.3:

$$\beta_S = [E(Y_i|T=1)_S - E(Y_i|T=0)_S] - \left[E(\boldsymbol{X}_i'\gamma|T=1)_S - E(\boldsymbol{X}_i'\gamma|T=0)_S\right]$$

From 3.23 and 3.3 it follows that

$$\beta_{S_1} \neq E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} \qquad (3.24)$$

Hence, for real RCTs, we can modify argument 3.22 and obtain the following sound deductive argument, using expression 3.24 as P_1:

$$\begin{array}{rl} P_1: & \beta_{S_1} \neq E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} \\ P_2: & E(Y_i|T=1)_{S_1} - E(Y_i|T=0)_{S_1} = 10,000€ \\ C: & \beta_{S_1} \neq 10,000€ \end{array} \qquad (3.25)$$

Hence, deductive logic tells us that $\beta_{S_1} \neq 10,000€$. Let's now go back to our argument of interest, argument [Job]. Since we already know that, in real scenarios, the conclusion of this argument is false, some readers might think it does not make sense to ask how much the premise of this argument supports its conclusion, and therefore there is no need for inductive logic in this context.

This challenge would require more space to be properly addressed – however, let me gesture at my response here. The underlying issue we need to address here is the nature of inductive support. Can a false proposition receive support from the evidence? The answer is affirmative, and our inductive logic should make sense of this. Vassend's verisimilitude framework for inductive inference (Vassend, 2020) provides some insight into this topic. As Vassend points out, "[i]t happens quite often that all the hypotheses under consideration are known to be false, because scientific hypothesizing very often (perhaps usually) takes place in highly idealized frameworks" (Vassend, 2020, p. 1363). However, we still believe that our evidence determines that some false hypotheses are closer to the truth than others. Our measure of inductive support, therefore, might be understood as a measure of the closeness to the truth (verisimilitude) of a conclusion given the truth of our premise. In Vassend's words, "the evidential measure should be calibrated to whatever the appropriate measure of verisimilitude is" (Vassend, 2020, p. 1367). In this scenario, the inductive support that our conclusion receives in argument [Job] might be understood as a measure of how close to the truth it is given the truth of the premises, i.e., of how close to the truth is the claim that $\beta_{S_1} = 10,000€$ despite being false, assuming that $E(Y_i|T = 1)_{S_1} - E(Y_i|T = 0)_{S_1} = 10,000€$. In its turn, verisimilitude might be understood in different ways – in chapter 2 I introduced the *content* approach, the *consequence* approach and the *likeness* approach to verisimilitude, for instance.

3.7 Conclusion

In this chapter I have illustrated the GFW and the HTI, by analysing a rule of causal inference in terms of the HTI, and I have explored two epistemic virtues of the GFW.

In sections 3.2 and 3.3 I have analysed a rule used for causal inferences in the context of comparative group studies (rule CIS). This analysis, together with the one provided in the previous chapter of rule EIS, should show the reader how the HTI can be applied to make sense of different rules of induction. The next chapter provides an additional example.

In section 3.3 I have also argued that analysing a rule of induction in terms of the HTI is a particularly clear approach. It is common to express rules of induction as slogans, which can then be articulated and applied in different ways. The HTI forces us to express our rule of interest precisely, in the form of a sentence schema, and this prevents confusions between similar but different rules. In its turn, this allows us to be explicit about the ideal warranting conditions for our rule of interest and to express them precisely, like the HTI requires.

In this chapter I have also highlighted two epistemic virtues of the GFW and the resulting version of the HTI. First, the ideal warranting conditions of a rule function as a methodological guide for researchers to ensure strong inferences. For example, in the context of comparative group studies, it is often the case that researchers are interested in arguments that can be understood from the perspective of rule CIS. Its ideal warranting conditions are $\mathrm{IWC}_{\mathrm{CIS}}$. There is disagreement in the literature about whether it is beneficial to randomly allocate subjects in a comparative group study. The HTI provides a clear argument in favour of randomization, since randomization brings us closer to $\mathrm{IWC}_{\mathrm{CIS}}$ thus ensuring stronger inferences. The HTI also provides a clear logical picture of why other methodological choices, like matching or blinding, are beneficial; similarly to randomization, these practices bring us closer to $\mathrm{IWC}_{\mathrm{CIS}}$ thus ensuring stronger inferences.

I have highlighted a second epistemic virtue of the GFW, namely, that the ideal warranting conditions of a rule function as a conceptual guide to the factors that are relevant in assessing an inductive inference's strength. This is the epistemic version of the logical point I made in §2.7. Argument [Job], for instance, can be understood from the perspective of rule CIS. Thus, IWC_{CIS} is a conceptual guide to the factors which are relevant in assessing the degree of support in argument [Job] and, thus, the strength of the corresponding inference; it tells us that these factors are those which are relevant for the truthlikeness of IWC_{CIS}([Job]), like the size of the study groups or the allocation method of participants.

Chapter 4

Bayesian Inductive Logics and the HTI

4.1 Introduction

In this chapter I explore Bayesian inductive logics from the perspective of the HTI. In doing so, I hope to illustrate three points.

First, that the HTI provides a unified framework to situate many open debates and disagreements about Bayesian inductive logics. In order to articulate a Bayesian inductive logic in the form of a sentence schema, as the HTI requires, one must take a stance with respect to several open debates. We must be explicit about our understanding of confirmation, whether (and how) we choose to quantify it, or what axioms of probability we prefer.[1] Without these commitments our position is not fully specified, and it cannot be articulated in the form of a sentence schema. The HTI provides a clear framework to understand which questions need to be answered, and how our answers impact the resulting logic.

Second, I will make explicit the central role that probability models play in any Bayesian inductive logic. In particular, I will bring to light the relation between the accuracy of our probability models and the accuracy of the information on inductive support that we obtain by using them in any Bayesian inductive logic. Any Bayesian rule must articulate the notion of confirmation in terms of probabilities. These probabilities are given by, or calculated from, a probability model of our system of interest. Hence, the ideal warranting conditions for any Bayesian rule will require that the probability model being used is a perfectly accurate model of our target system. Under these ideal conditions, the information on inductive support that our Bayesian rule provides is true; as our probability model becomes less accurate, the information on confirmation provided by our Bayesian rule becomes less truthlike. This is the general understanding of Bayesian inductive logics that the HTI provides.

Last, I will show that understanding Bayesian inductive logics from the perspective of the HTI

[1]The terms "confirmation" and "inductive support" are used interchangeably in the context of inductive logic. For instance, when introducing the aim of inductive logic, Hájek and Joyce tell us that "[t]he central task is to understand what it means to say that datum E *confirms or supports* a hypothesis H when E does not logically entail H" (Hájek and Joyce, 2008, p. 115, my emphasis). I will also use these terms interchangeably in this chapter, since "confirmation" is a very popular term in the Bayesian literature.

has some epistemic consequences too. Recall, from chapter 3, that the ideal warranting conditions of a given rule can function as a methodological and as a conceptual guide for researchers. The ideal warranting conditions for any Bayesian rule require our probability model to be perfectly accurate. In the current context, these ideal warranting conditions function as a methodological guide in the following sense: they tell us that we should maximize the accuracy of our models, since more accurate models will yield more truthlike information on inductive support, which will better inform our inferences. This is because the more accurate the information on inductive support that we have, the less guesswork is involved in correcting for deviations due to the model's inaccuracies. These ideal warranting conditions also function as a conceptual guide to the factors that are relevant in assessing the strength of our inferences: these will be the factors which are relevant in assessing the truthlikeness of the ideal warranting conditions for our Bayesian rule, hence, the factors which are relevant for the accuracy of our probability models. These epistemic consequences are in line with some recent suggestions in the literature on Bayesian epistemology. Morey et al. (2013) and Gelman and Shalizi (2013), for example, urge that philosophical accounts of Bayesianism should account for the practice of model-checking (that is, checking one's own models). The need for model-checking follows from looking at Bayesianism from the perspective of the HTI: we must check our models in order to assess their accuracy, so that we can assess the strength of our inferences. Relatedly, Sprenger and Hartmann (2019, pp. 311–26) suggest a particular understanding of Bayesianism as a form of model-based reasoning, consistent with the epistemic consequences just introduced. Thus, the HTI offers a logical underpinning for these epistemic positions. Furthermore, by making the ideal warranting conditions of Bayesian rules explicit, the HTI can help us develop better bridge principles to understand the connection between Bayesian inductive logics and Bayesian epistemologies. In this chapter I illustrate how we can go about this, by showing how the HTI allows us to generalize a specific bridge principle for Bayesian inductive logics suggested by Fitelson (2006).

Here is how I will proceed. In section 4.2 I introduce the distinction between Bayesian inductive logics and Bayesian epistemologies to prevent misunderstandings. In section 4.3 I articulate a specific Bayesian rule. In doing so, I will situate several open debates regarding Bayesian inductive logic in a common framework. In section 4.4 I will expose the ideal warranting conditions for our Bayesian rule. This will illuminate the connection between model accuracy and inductive support for this rule. Section 4.5 generalizes the argument from section 4.4, showing that, for any Bayesian rule, the truthlikeness of the information on confirmation the rule provides depends on the accuracy of the probability model being used. In section 4.6 I explore some epistemic consequences of understanding Bayesian inductive logics from the perspective of the HTI. Section 4.7 concludes.

4.2 Bayesian Inductive Logics and Bayesian Epistemologies

The term "Bayesianism" is used to refer to a diversity of views. At the most general level, we should distinguish between Bayesian inductive logics and Bayesian epistemologies, different types of accounts that often fall under the common term "Bayesianism". In this section I clarify this distinction to prevent confusions.

As introduced in chapter 1, by "inductive logic" I mean an account of the relations of support between *propositions* in an ampliative argument (Hawthorne, 2018; Norton, 2021c). In contrast, epistemology is concerned with rational/justified *belief*. Logic and epistemology are distinct but closely related, and, as a result, they are often conflated or confused. In fact, the term "logic" is sometimes used to refer to theories of rational belief or belief revision. This is problematic.

Harman (1984, 2002, 2009) diagnosed this problem for the case of deductive logic with great clarity. Harman contrasts deductive logic with what he calls "theories of reasoning". He argues that both endeavours are often mistakenly conflated. The subject matter of theories of reasoning are the psychological events or processes that constitute reasoning, while the subject matter of deductive logic are relations between propositions (Harman, 2009). Thus, "logical principles are not directly rules of *belief revision*. They are not particularly about belief at all" (Harman, 1984, p. 107, emphasis in original). Once we realise this, we see that the gap between logic and reasoning must be bridged. The challenge of bridging this gap is known as "Harman's challenge" (Steinberger, 2022). Thus, if we think that logic is normative for thought in one way or another, we must articulate this relation in the form of a "bridge principle". The principle of closure under logical entailment is a famous bridge principle for deductive logic. This principle tells us that if a subject S knows that p, and p entails q, then S knows that q. However, articulating and defending a good bridge principle for deductive logic is harder than it seems, as MacFarlane (2004) shows.[2]

An analogous challenge exists in inductive logic. A Bayesian inductive logic must articulate confirmation in terms of *objective* probabilities, since confirmation (as used in the context of inductive *logic*) is a relation of support between propositions and, therefore, between possible states of the world.[3] Thus, the probabilities involved in representing this relation of support must be themselves representing something about the world, and objective probabilities do just that:[4]

> [O]bjective chances are commonly taken to make empirical statements: their values are constituted by patterns of events and processes in the actual world (Hoefer 2007, 549), such as the setup of an experiment or the physical composition of the coin we toss. This is true of frequencies, propensities and best-system chances alike. (Sprenger and Hartmann, 2019, p. 315)

However, we want the information on confirmation to bear on our epistemic attitudes too. For instance, *ceteris paribus*, the more a a piece of evidence E confirms a hypothesis H, the higher our credence in H should be. And credences are the subject matter of Bayesian epistemologies. Thus, we must find ways to bridge Bayesian inductive logics and Bayesian epistemologies. We can see this challenge as an analogous version of Harman's challenge for inductive logics. Chance-credence coordination principles (like Lewis' Principal Principle (Lewis, 1980), or the Principle

[2]See (Hawthorne, 2004) for a defence of deductive closure principles in epistemology.

[3]This is in line, for example, with Sober's take on Bayesianism, a conclusion he reaches through other means: "[w]hen prior probabilities can be defended empirically, and the values assigned to a hypothesis' likelihood and to the likelihood of its negation are also empirically defensible, you should be a Bayesian" (Sober, 2008, p. 32).

[4]Articulating confirmation in terms of objective probabilities is just a minimal condition for an inductive logic to be Bayesian. There are other logics that also articulate confirmation in terms of objective probabilities and which are not Bayesian. For instance, "[t]here is a view, or family of views, called *likelihoodism* that maintains that the inductive logician or statistician should only be concerned with whether the evidence provides *increased* or *decreased support* for one hypothesis over another, and only in cases where this evaluation is based on the ratios of *completely objective* likelihoods" (Hawthorne, 2018, emphasis in original). Thus, likelihoodists argue that confirmation is an intrinsically comparative notion, so likelihoodism can only tell us whether E confirms H_1 to a greater or lesser extent than H_2. In contrast, Bayesian inductive logics tell us whether E confirms H simpliciter and, maybe, to which degree.

of Direct Inference (Reichenbach, 1949)) are classic attempts at bridging this gap.[5] But one can also attempt to bridge confirmation and other epistemic concepts; Fitelson (2006), for example, defends a bridge principle for Bayesian inductive logics bridging confirmation and incremental evidential support (an epistemic concept). At the end of the chapter I will explore some epistemic consequences of understanding Bayesian inductive logics from the perspective of the HTI, and I will briefly talk about the necessary bridge principles.

In this section I have introduced the distinction between Bayesian inductive logics and Bayesian epistemologies. The subject matter of Bayesian inductive logics are relations of support between propositions, which are represented in terms of objective probabilities. In contrast, the subject matter of Bayesian epistemologies are credences or degrees-of-belief, and their relations. These credences are represented in terms of subjective probabilities. Logic and epistemology are closely related, but distinct. In this chapter I will articulate a Bayesian inductive logic in terms of the HTI (§4.3 to §4.5) and I will explore some implications of this analysis for Bayesian epistemology (§4.6).

4.3 Articulating a Bayesian Rule

In this section I will articulate one rule corresponding to a Bayesian inductive logic; in the next section I will articulate the ideal warranting conditions for this rule.

With this analysis I hope to illustrate two points. First, that the HTI provides a unified framework within which many open debates and disagreements about Bayesian inductive logic can be situated. Thinking about Bayesian inductive logics from the perspective of the HTI will help us understand what these debates are about, what part of the Bayesian apparatus is at stake and what implications they have for our logic. Second, and crucially, in expressing Bayesian inductive logics in terms of the HTI, the role that the accuracy of our probability models plays in our Bayesian accounts of confirmation will become apparent.

In order to articulate a Bayesian inductive logic we will need to take a stance on several open debates, so that our view is fully specified and can be expressed in terms of the HTI. This will hopefully be an illuminating process, since it will situate each of these debates in a common framework. However, it is important to remember that the HTI does *not* tell us how to make these commitments, and neither should we expect it to; the HTI provides a framework to articulate our inductive logic once all the necessary commitments have been made at the outset. There are many other Bayesian inductive logics available, which will differ from the one articulated here in one or more of these underlying commitments and, therefore, will be expressed by a different rule (see §4.5 for one example). This is no challenge to the HTI. We can use the HTI as a framework to articulate our favourite Bayesian inductive logic. This section and the next provide just one example of what that process looks like.

In this section I articulate my rule of interest, which tells us that the degree of support (or confirmation) that a piece of evidence E yields a hypothesis H is determined by a three-place confir-

[5] In fact, Carnap (1950/1962, p. 201) introduces a closure principle for deductive logic and a confirmation-credence coordination principle for inductive logic side-by-side, as equivalent bridge principles.

mation function $\mathfrak{c}(H, E, \mathcal{M})$, where $\mathcal{M}$ is a probability model. This rule can be expressed in the form of the following sentence schema:

BAY_1: E confirms H to degree $\mathfrak{c}(H, E, \mathcal{M})$

with the corresponding side condition, which I will provide shortly.

I call this rule BAY_1 because it is just one of multiple Bayesian rules $\{BAY_1, BAY_2, \ldots, BAY_n\}$ which correspond to multiple Bayesian inductive logics. Each rule will differ either with regards to its template or with regards to its side condition, differences that will reflect diverse underlying commitments.

We must now articulate the side condition of BAY_1. Recall that a rule of induction is a sentence schema, consisting of a template and a side condition (see chapter 2). The template is a syntactic string containing significant words or symbols and also placeholders. The side condition must specify how to fill in those placeholders, and it may also specify how to understand the significant words or symbols. The placeholders in BAY_1 are E, H and $\mathcal{M}$. In addition to these, BAY_1 contains the significant word "confirms" and the symbol "$\mathfrak{c}$", which the side condition must explicate.[6] In previous chapters, all the rules I have used as examples have posited an entailment relation, which requires no clarification. However, Bayesian rules will posit a weaker confirmation relation, which can be understood in several ways. Accordingly, the side condition of a Bayesian rule must clarify how we should make sense of confirmation in probabilistic terms. Similarly, symbol $\mathfrak{c}$ refers to a measure of confirmation, and our side condition must specify what that measure is. In sum, the side condition of BAY_1 must indicate how to understand the term "confirms" and the symbol "$\mathfrak{c}$", and how to fill in placeholders E, H and $\mathcal{M}$. I now turn to this task.

4.3.1 Explicating confirmation and $\mathfrak{c}$

Let's start by explicating confirmation. There are several ways one can understand confirmation. Disagreements about confirmation may arise, first, at the metaphysical level; what is the nature of confirmation? This matter, of course, is far from settled. Carnap (1950/1962) thought of confirmation as a relation between the ranges of propositions. The range of a proposition, $\mathfrak{R}(\cdot)$, is the set of possible worlds where this proposition is true.[7] Thus, according to Carnap, to say that E confirms H (to a certain degree) is to say something about the relation between $\mathfrak{R}(E)$ and $\mathfrak{R}(H)$. This fundamental understanding of inductive support is still popular nowadays, and is adopted, for example, by Hawthorne (2018). While there isn't much literature on the metaphysics of confirmation, some contemporary positions seem to challenge the Carnapian view. Vassend (2020), for example, takes inductive support to capture the verisimilitude of the conclusion of an inductive argument.[8] Thus, from this perspective, to say that E confirms H (to a certain degree) is to say

[6]Notice that $\mathfrak{c}$ is not defined as a placeholder here but as a significant symbol. Thus, it will always take the same value, in all instances of the rule. In contrast, placeholders may be filled in differently for each instance of the rule. As a consequence, if one wishes to, $\mathfrak{c}$ can be substituted by its definition in the rule's template.

[7]Although Carnap uses the term "state-descriptions" instead of "possible worlds".

[8]Vassend uses the concept of "evidential favouring" instead of "inductive support". However, these two concepts are closely related and can be bridged. Fitelson and Sober suggest the following principle: "E favors H_1 over H_2 iff E supports H_1 more strongly than E supports H_2" (Fitelson and Sober, 2011, p. 667). Despite being an intuitive principle some likelihoodsits might disagree with it, as Fitelson and Sober acknowledge. See (Fitelson, 2007; Fitelson & Sober,

something about the verisimilitude of H given the truth of E. Of course, the nature of "verisimilitude" is itself a matter of debate, as discussed in chapter 2. For example, if we favour a "content approach" to verisimilitude, then the verisimilitude of H depends on the amount of true content of H and, accordingly, under Vassend's view, the inductive support that E yields H would be indicative of the amount of true content of H assuming the truth of E.

However, most work on confirmation skips the metaphysical debate altogether and focuses on operationalising this concept; disagreements arise at this level too. Operationalising the notion of confirmation means defining it in such a way that it can be applied. Following Hempel (1945a, 1945b) and Carnap (1950/1962) we can distinguish three types of operationalisations of confirmation: qualitative (or classificatory), comparative (or relational) and quantitative. Hempel thought that operationalising confirmation had to start at the qualitative stage, and then move on to the comparative stage, before we could reach the end goal: a quantitative measure of confirmation. While this is a useful taxonomy, we don't have to accept the hierarchy between these types of operationalisations. Likelihoodists think that confirmation is an intrinsically comparative notion. We may also think that qualitative operationalisations of confirmation are more objective than quantitative operationalisations of confirmation (Fitelson, 2006, p. 506). Crucially, how we operationalise confirmation will depend on what we take it to capture and how we want it to behave.

Regarding qualitative operationalisations of confirmation, there is almost universal agreement among Bayesians that E confirms H iff $P(H|E) > P(H)$. This is a qualitative operationalisation of confirmation because, by itself, it only allows us to determine whether H is confirmed by E or not. Carnap (1950/1962) called this fundamental understanding of confirmation "confirmation as increase in firmness". We can contrast this with Carnap's alternative notion of "confirmation as firmness", an intrinsically quantitative notion of confirmation according to which $P(H|E)$ is itself a measure of how much E confirms H. Using Hájek's terminology (2008), "confirmation as increase in firmness" is an incremental notion of confirmation, while "confirmation as firmness" is an absolute notion of confirmation. Fitelson (2006) argues that confirmation as increase in firmness is better than confirmation as firmness, in that the former, but not the latter, satisfies the three Carnapian desiderata of analyticity, logicality and applicability (when suitably understood).[9] Of course, our agreement with Fitelson hinges on whether we take these to be suitable desiderata.

Regarding comparative operationalisations of confirmation, the consensus among Bayesians is that comparative notions of confirmation are derived concepts, defined in terms of a primitive, non-comparative confirmation concept (Fitelson, 2007, p. 474). In this way, most Bayesians think that the comparative statement "E favours H_1 over H_2" is equivalent to "$\mathfrak{c}(H_1, E) > \mathfrak{c}(H_2, E)$", where $\mathfrak{c}$ is some confirmation function that captures the degree to which a piece of evidence supports a hypothesis *simpliciter*. Carnap expressed this point neatly: "the problem of finding an adequate explicatum in non-quantitative terms for the comparative concept of confirmation is at present unsolved" (Carnap, 1953, p. 312). Most Bayesians don't think that this problem can be solved; thus, they think that any comparative notion of confirmation is derived from a fundamental quantitative notion of confirmation. Likelihoodists disagree with Bayesians on this point (Fitelson, 2007). According to likelihoodists, the notion of confirmation is intrinsically comparative and is captured by the Law of Likelihood: E supports H_1 over H_2 precisely when

2011) for a more extended discussion of this principle and its consequences.

[9] Fitelson modifies Carnap's deisderatum of applicability.

$P(E|H_1) > P(E|H_2)$. Likelihoodists do not talk of E supporting H_1 simpliciter. The Law of Likelihood is a comparative operationalisation of confirmation because, by itself, it only allows us to tell whether a particular hypothesis is more or less confirmed by the evidence than an alternative hypothesis.[10]

Last, regarding quantitative operationalisations of confirmation, there is little consensus among Bayesians. As a result, there are plenty of confirmation measures available for Bayesians to choose from (Kyburg, 1983; Roche & Shogenji, 2014). The difference measure, d, the log-ratio measure, r, and the log-likelihood measure, l, are three popular measures of confirmation in the literature (Fitelson, 1999):

$$\begin{aligned} d(H,E,\mathcal{M}) &= P_{\mathcal{M}}(H|E) - P_{\mathcal{M}}(H) \\ r(H,E,\mathcal{M}) &= log\left[\frac{P_{\mathcal{M}}(H|E)}{P_{\mathcal{M}}(H)}\right] \\ l(H,E,\mathcal{M}) &= log\left[\frac{P_{\mathcal{M}}(E|H)}{P_{\mathcal{M}}(E|\neg H)}\right] \end{aligned}$$

Notice that I have made confirmation measures $\mathfrak{c}$ three-place functions, $\mathfrak{c}(H,E,\mathcal{M})$, while they are normally expressed as two-place functions, $\mathfrak{c}(H,E)$. In this approach I am following Fitelson (2006). As it will become evident shortly, there are multiple ways that a system can be probabilistically modelled, and, correspondingly, there are different ways one can assign probabilities for the outcomes and events of that system. In this way, $P_{\mathcal{M}}$ stands for an assignment of probabilities according to probability model $\mathcal{M}$.[11] Thus, the value of any measure $\mathfrak{c}$ will depend on our choice of probability model $\mathcal{M}$, and I choose to make that dependence explicit.

Measures d, r and l can all be understood as quantitative measures of "confirmation as increase in firmness". This is so because each of these measures satisfies the following condition called "sensitivity to increase in firmness":

$$\mathfrak{c}(H,E,\mathcal{M}) \begin{cases} > n, & \text{if } P_{\mathcal{M}}(H|E) > P_{\mathcal{M}}(H) \\ = n, & \text{if } P_{\mathcal{M}}(H|E) = P_{\mathcal{M}}(H) \\ < n, & \text{if } P_{\mathcal{M}}(H|E) < P_{\mathcal{M}}(H) \end{cases}$$

Measures d, r and l are designed so that $n = 0$, which is sometimes adopted as a convention. I will adopt this convention here too.

If, instead of understanding confirmation as increase in firmness, we choose to understand confirmation as firmness, then "Bayesian confirmation theory is little more than the examination of [the] properties" of the posterior probability function (Howson, 2000, p. 179). In this case, we should simply use the posterior, $P_{\mathcal{M}}(H|E)$, as a measure of confirmation (e.g. Hawthorne, 2018).

[10] Some likelihoodists do endorse that the likelihood itself, $P(E|H)$, is a measure of the degree to which E confirms H. In that case, the boundary between likelihoodism and Bayesianism starts to blur. These distinctions are not crucial for my arguments in this chapter, which apply to any inductive logic where confirmation is understood in terms of probabilities.

[11] See also (Sprenger & Hartmann, 2019, pp. 311–26) for this notation.

For the purposes of articulating rule BAY_1 let us commit to understanding confirmation as increase in firmness, and to the idea that this increase in firmness can be quantitatively measured. Different commitments would just result in different rules. Which confirmation measure(s) we favour will depend on which adequacy conditions we impose on them. Adequacy conditions capture our core intuitions about what these measures are measuring and how they should behave. Thus, there is no such thing as a "correct" measure of confirmation, but only "adequate" measures, which satisfy our preferred adequacy conditions.[12]

Accordingly, it is crucial that we are explicit in what we take to be the appropriate adequacy conditions for measures of confirmation. Fitelson (2006, p. 506) suggests the following two adequacy conditions:

(AC1)

$$\mathfrak{c}(H, E, \mathcal{M}) \begin{cases} \text{Maximal if } H \text{ entails } E \\ > 0, \ \text{ if } \ P_{\mathcal{M}}(H|E) > P_{\mathcal{M}}(H) \\ = 0, \ \text{ if } \ P_{\mathcal{M}}(H|E) = P_{\mathcal{M}}(H) \\ < 0, \ \text{ if } \ P_{\mathcal{M}}(H|E) < P_{\mathcal{M}}(H) \\ \text{Minimal if } H \text{ entails } \neg E \end{cases}$$

(AC2) If $P_{\mathcal{M}}(H|E_1) > P_{\mathcal{M}}(H|E_2)$ then $\mathfrak{c}(H, E_1, \mathcal{M}) > \mathfrak{c}(H, E_2, \mathcal{M})$

The first condition, (AC1), is jointly expressing two separate commitments:

- Logicality: $\mathfrak{c}(H, E, \mathcal{M})$ should be a quantitative generalization of entailment. Thus, it should be maximal when H entails E, and minimal when H entails $\neg E$.
- Sensitivity to increase in firmness: there should be a value n such that $\mathfrak{c}(H, E, \mathcal{M})$ is greater than n when there is increase in firmness, equal to n when there is no increase nor decrease in firmness, and smaller than n when there is decrease in firmness (with the added convention that $n = 0$).

Fitelson favours measure l since it satisfies (AC1) and (AC2), which he takes to be appropriate adequacy conditions. In fact, l (and its ordinal equivalents) is the only historically proposed measure of confirmation satisfying (AC1) and (AC2) (Fitelson, 2006, p. 506). Of course, one may object that (AC1) and (AC2) are not appropriate adequacy conditions. Glass and McCartney (2015, pp. 62–63), for example, reject logicality as an adequacy condition, and suggest some alternative adequacy conditions instead. Zalabardo (2009) suggests still a different set of adequacy conditions. However, as long as we take (AC1) and (AC2) to be appropriate (i.e., necessary and sufficient) then we must also accept that l is an adequate measure of confirmation. For the purposes of articulating rule BAY_1, let us accept Fitelson's adequacy conditions and, as a result, let us use l as a measure of confirmation.[13]

[12] One might defend that there is a privileged set of adequacy conditions that singles out the one true measure of confirmation (Milne, 1996). This position is known as confirmational monism, and it is problematic. Sprenger and Hartmann (2019, p. 60), for instance, lay out two plausible adequacy conditions that no measure of confirmation jointly satisfies. Thus, there seem to be different explications of confirmation available, that may be captured by different measures, a position known as confirmational pluralism (e.g. Fitelson, 1999, 2001; Sprenger & Hartmann, 2019).

[13] I pick measure l due to its popularity; however, notice that this choice is irrelevant for the arguments I present in this chapter. The reader might insert her favourite set of adequacy conditions, and her favourite measure of confirmation, and the arguments in this chapter remain unchanged. This will become evident in §4.5.

Recall that the side condition for rule BAY_1 must explicate the notion of "confirmation". We are now ready to do this. Since, for the purposes of articulating this rule, I committed to understanding confirmation as increase in firmness and I accepted Fitelson's adequacy conditions, the side condition for BAY_1 will tell us that "confirmation is increase in firmness, and an adequate measure of increase in firmness must satisfy (AC1) and (AC2)".

Our side condition must also explicate $\mathfrak{c}$. Remember that $\mathfrak{c}$ is not a placeholder here, but a significant symbol denoting an adequate measure of confirmation. Our choice of $\mathfrak{c}$ must be consistent with our understanding of confirmation; that is, our function $\mathfrak{c}$ must be a measure of whatever it is that we understand as confirmation, which we have just explicated. Since, for rule BAY_1, "confirmation is increase in firmness, and an adequate measure of increase in firmness must satisfy (AC1) and (AC2)", and since measure l satisfies (AC1) and (AC2), then we can make $\mathfrak{c} = l$ (the log-likelihood measure).

4.3.2 How to fill in the placeholders

Let me now turn to the placeholders in BAY_1: E, H and $\mathcal{M}$. The side condition of BAY_1 must specify how to fill in these placeholders. Placeholders E and H must be filled in with the premise P and the conclusion C of our argument of interest, respectively. Suppose we are interested in argument [Die]. In that case, E="A fair six-sided die has landed on an even number" and H="The die has landed on 2".[14]

P :	A fair six-sided die has landed on an even number	[Die]
C :	The die has landed on 2	

Placeholder $\mathcal{M}$ must be filled in with the identifier of a specific probability model of our system. A probability model, probability space, or probability triple, is a mathematical construct consisting of three elements, $(\Omega, \mathcal{F}, P)$, which must satisfy a given set of axioms of probability.

Let me discuss each of these ingredients of $\mathcal{M}$ while putting together one possible model of our system of interest, the throw of a fair six-sided die, which I will call $\mathcal{M}_{1,[\text{Die}]}$. I will call this model $\mathcal{M}_{1,[\text{Die}]}$ since, as it will become evident, our system may be modelled in many different ways, thus obtaining alternative models $\{\mathcal{M}_{2,[\text{Die}]}, \mathcal{M}_{3,[\text{Die}]}, \ldots, \mathcal{M}_{n,[\text{Die}]}\}$.

Since our side condition must specify how to fill in placeholder $\mathcal{M}$, it must provide a set of axioms of probability that our probability model $\mathcal{M}$ must satisfy. Kolmogorov's (1933/1950) axioms are the most popular, but there are alternative options available. Hawthorne (2018), for instance, provides an axiomatization of probability that takes conditional probability as basic, in contrast with Kolmogorov's axioms, which take unconditional probability as basic. Different axiomatizations of probability will result in different side conditions for our template (since the instructions to fill in placeholder $\mathcal{M}$ will differ) and, as a result, in different Bayesian rules. Given the popularity of

[14] One could argue that some of the information included in E is not part of the evidence, but should be considered "background knowledge" K. For instance, we could restrict the evidence to E="The die has landed on an even number", and make K="The die is fair and has six sides". This approach is also fine and would not change the analysis I provide here beyond the fact that all the probabilities would be conditioned on K.

Kolmogorov's axioms, for the purposes of articulating rule BAY_1 I will adopt this axiomatization of probability.

We can now define the elements of our probability model $\mathcal{M}$, which must satisfy Kolmogorov's axioms:

- The sample space Ω is a non-empty set containing all possible outcomes of our system. An outcome w is the result of a single execution of the system.
- The event space $\mathcal{F}$ is a σ-algebra on Ω, containing all events S we would like to consider.[15] An event S is a set of outcomes w. Thus, $F \subseteq 2^{\Omega}$.[16] An event is said to "have happened" if the outcome of the system is an element of that event.
- The probability function P assigns every event in $\mathcal{F}$ a number between 0 and 1, $P : \mathcal{F} \to [0, 1]$

With this, we can now construct $\mathcal{M}_{1,[\text{Die}]} = \left(\Omega_{1,[\text{Die}]}, \mathcal{F}_{1,[\text{Die}]}, P_{1,[\text{Die}]}\right)$, our model for the throw of a fair six-sided die.

First, we must define the sample space $\Omega_{1,[\text{Die}]}$. The outcomes in any sample space must be mutually exclusive and collectively exhaustive. An obvious choice of sample space for the throw of a six-sided die is the set $\{1, 2, 3, 4, 5, 6\}$, where every element stands for the number of pips facing up, but this is not our only choice. The results of two different executions of a system only count as different outcomes if the differences between them matter for our analysis and purposes. If we are interested in the number of pips in the upward-facing face of the die, then set $\{1, 2, 3, 4, 5, 6\}$ is indeed an adequate choice of sample space. If, however, we are playing a simple game where all that matters is whether the die lands on an even number or an odd number, then $\{E, O\}$ (where E stands for "even" and O for "odd") would be an appropriate choice of sample space. When playing this simple game, the die landing on 2 and the die landing on 4 count as the same outcome, since they do not differ in any way that matters for our purposes. In other words, our sample space must have the right granularity, and this is determined by our purposes. In order to understand argument [Die], $\{1, 2, 3, 4, 5, 6\}$ is an appropriate choice of sample space, since it allows us to capture the differences that are relevant to us in the results of executions of the system. In contrast, $\{E, O\}$ is an inappropriate choice of sample space (it is too coarse) since it does not allow us to capture the outcome "the die has landed on 2". Similarly, we can come up with sample spaces that are too fine given our purposes. Take the set $\{(1, D), (1, N), (2, D), (2, N), \ldots, (6, D), (6, N)\}$, where $(1, D)$ means that the die has landed on 1 and it is daytime, and $(1, N)$ means that the die has landed on 1 and it is nighttime. This set would also be an inappropriate sample space, since the differences between $(1, D)$ and $(1, N)$ are irrelevant given our purposes. In this case, the sample space would be too fine. Thus,

$$\Omega_{1,[\text{Die}]} = \{1, 2, 3, 4, 5, 6\}$$

However, we should keep in mind that other sample spaces are also available to us and may be adequate in other circumstances.

[15] A σ-algebra on a set S is a nonempty collection of subsets of S closed under complement, countable unions and countable intersections.

[16] Note that 2^{Ω} is the powerset of Ω, that is, the set of all subsets of Ω including the empty set and Ω itself.

Then we must define our event space $\mathcal{F}_{1,[\text{Die}]}$. $\mathcal{F}_{1,[\text{Die}]}$ must be a σ-algebra on $\Omega_{1,[\text{Die}]}$. Of course, different choices of Ω may result in different $\mathcal{F}$s, but we may also obtain different $\mathcal{F}$s for a fixed Ω. It is up to us to pick a σ-algebra on $\Omega_{1,[\text{Die}]}$ that contains all the events we would like to consider. For instance, the set $\{\emptyset, \{3,4\}, \{1,2,5,6\}, \{1,2,3,4,5,6\}\}$ is a σ-algebra on $\Omega_{1,[\text{Die}]}$ and, therefore, it is one possible event space available to us. However, this event space is not useful to us. Given $\Omega_{1,[\text{Die}]}$, and given our purposes and desired analysis, our event space must at least contain the event $\{2,4,6\}$ (the die has landed on an even number) and $\{2\}$ (the die has landed on 2). There are multiple σ-algebras on $\Omega_{1,[\text{Die}]}$ containing these two events. For simplicity, I will pick

$$\mathcal{F}_{1,[\text{Die}]} = 2^{\Omega_{1,[\text{Die}]}}$$

which is the greatest σ-algebra on $\Omega_{1,[\text{Die}]}$ and, therefore, it contains all possible events.

Last, we must define our probability function $P_{1,[\text{Die}]}$. P is a set function returning an event's probability, $P : \mathcal{F} \longrightarrow [0,1]$. For the throw of a *fair six-sided* die the probability mass function $p_{1,[\text{Die}]} : \Omega_{1,[\text{Die}]} \longrightarrow [0,1]$ is the constant function $p_{1,[\text{Die}]}(w) = 1/6$, since this is what it means for a six-sided die to be fair. As a result, our probability function P is

$$P_{1,[\text{Die}]}(S) = \sum_{w \in S} p_{1,[\text{Die}]}(w) = \frac{n(S)}{6}$$

where $n(S)$ is the cardinality of set S, i.e., the number of elements in set S. In this way, for event $\{2,3,4\}$, $P_{1,[\text{Die}]}(\{2,3,4\}) = n(\{2,3,4\})/6 = 3/6 = 0.5$.

We have now fully specified how to fill in placeholder $\mathcal{M}$: $\mathcal{M}$ is a probability model of our system, consisting of three elements, $(\Omega, \mathcal{F}, P)$, where Ω is the sample space, $\mathcal{F}$ the event space, and P a probability measure, $P : \mathcal{F} \longrightarrow [0,1]$. These elements must satisfy Kolmogorov's axioms.

4.3.3 Summing up

We can now fully articulate rule BAY_1, in the form of a template with its corresponding side condition:

BAY_1: E confirms H to degree $\mathfrak{c}(H, E, \mathcal{M})$

where

(1) Confirmation is increase in firmness, and an adequate measure of increase in firmness must satisfy (AC1) and (AC2);

(2) $\mathfrak{c} = l$ (the log-likelihood measure);

(3) E is the premise of our argument of interest;

(4) H is the conclusion of our argument of interest;

(5) $\mathcal{M}$ is a probability model of our system, consisting of three elements, $(\Omega, \mathcal{F}, P)$, where Ω is the sample space, $\mathcal{F}$ the event space, and P a probability measure, $P : \mathcal{F} \to [0, 1]$. These elements must satisfy Kolmogorov's axioms.

Let's explore how we would go about generating an instance of rule BAY_1. The template in BAY_1 has three placeholders: E, H and $\mathcal{M}$. Placeholders E and H can always be filled in just by looking at our argument of interest A; placeholder $\mathcal{M}$, however, is filled in by independently providing a probability model. Going back to the notation introduced in chapter 2, $\text{R}\,(\text{A}, \Phi_1, \ldots, \Phi_n)$ is the instance of rule R evaluated in context $(\text{A}, \Phi_1, \ldots, \Phi_n)$, where A refers to our argument of interest, and $\Phi_1, \ldots, \Phi_n$ are all placeholders in R which are *not* filled in with information about our argument of interest. Rule BAY_1 has one placeholder which is not filled in with information about our argument of interest: placeholder $\mathcal{M}$. Thus, in order to generate an instance of the rule BAY_1 we need information both contained in our argument of interest A and in a given probability model $\mathcal{M}$, obtaining the sentence $\text{BAY}_1\,(A, \mathcal{M})$.

This means that an instance of rule BAY_1, $\text{BAY}_1\,(A, \mathcal{M})$, can only be formed whenever we can provide a probability model $\mathcal{M}$ of our target system. Whenever our target system cannot be probabilistically modelled, placeholder $\mathcal{M}$ cannot be filled and our argument of interest cannot be understood in terms of rule BAY_1. Recall the debate between Norton and Earman about Norton's dome, introduced in chapter 1; Norton (2010) argues that this system cannot be described in probabilistic terms, since the physics of the system do not provide physical chances for its possible futures, while Earman (2020) disagrees. This is a debate about whether we can provide a probability model $\mathcal{M}$ for Norton's dome.

It is now interesting to see what an instance of our rule would look like. Let's focus on argument [Die]:

P :	A fair six-sided die has landed on an even number	[Die]
C :	The die has landed on 2	

I pick model $\mathcal{M}_{1,\,[\text{Die}]}$, developed above, as the probability model of our system. For clarity, let me refer to the premise P of our argument as the evidence in argument [Die], $E_{[\text{Die}]}$. Similarly, let me refer to the conclusion C of our argument as the hypothesis in argument [Die], $H_{[\text{Die}]}$. $E_{[\text{Die}]}$ corresponds to event $\{2, 4, 6\}$ in model $\mathcal{M}_{1,\,[\text{Die}]}$; $H_{[\text{Die}]}$ corresponds to event $\{2\}$ in model $\mathcal{M}_{1,\,[\text{Die}]}$; $\neg H_{[\text{Die}]}$ corresponds to event $\{1, 3, 4, 5, 6\}$ in model $\mathcal{M}_{1,\,[\text{Die}]}$. Then:

$$\text{BAY}_1\left([\text{Die}], \mathcal{M}_{1,\,[\text{Die}]}\right)\text{: } E_{[\text{Die}]} \text{ confirms } H_{[\text{Die}]} \text{ to degree } l(H_{[\text{Die}]}, E_{[\text{Die}]}, \mathcal{M}_{1,\,[\text{Die}]})$$

We can calculate the value of our measure of confirmation:

$$
\begin{aligned}
l(H_{[\text{Die}]}, E_{[\text{Die}]}, \mathcal{M}_{1,[\text{Die}]}) &= log\left[\frac{P_{1,[\text{Die}]}(E_{[\text{Die}]}|H_{[\text{Die}]})}{P_{1,[\text{Die}]}(E_{[\text{Die}]}|\neg H_{[\text{Die}]})}\right] \\
&= log\left[\frac{P_{1,[\text{Die}]}(\{2,4,6\}|\{2\})}{P_{1,[\text{Die}]}(\{2,4,6\}|\{1,3,4,5,6\})}\right] \\
&= log\left[\frac{\dfrac{P_{1,[\text{Die}]}(\{2,4,6\}\cap\{2\})}{P_{1,[\text{Die}]}(\{2\})}}{\dfrac{P_{1,[\text{Die}]}(\{2,4,6\}\cap\{1,3,4,5,6\})}{P_{1,[\text{Die}]}(\{1,3,4,5,6\})}}\right] \\
&= log\left[\frac{\dfrac{P_{1,[\text{Die}]}(\{2\})}{P_{1,[\text{Die}]}(\{2\})}}{\dfrac{P_{1,[\text{Die}]}(\{4,6\})}{P_{1,[\text{Die}]}(\{1,3,4,5,6\})}}\right] = log\left[\frac{\dfrac{n(\{2\})/6}{n(\{2\})/6}}{\dfrac{n(\{4,6\})/6}{n(\{1,3,4,5,6\})/6}}\right] \\
&= log\left[\frac{\dfrac{1/6}{1/6}}{\dfrac{2/6}{5/6}}\right] = log(5/2) \approx 0.4
\end{aligned}
$$

Thus,

$\text{BAY}_1\left([\text{Die}], \mathcal{M}_{1,[\text{Die}]}\right)$: $E_{[\text{Die}]}$ confirms $H_{[\text{Die}]}$ to degree 0.4

Recall that instances of rules of induction inform us about relations of inductive support. As expected, $\text{BAY}_1\left([\text{Die}], \mathcal{M}_{1,[\text{Die}]}\right)$ is informing us about the relation of inductive support in argument [Die]. In particular, it tells us that the degree of support in that argument is 0.4. In order to understand what this number means, we must look into the side condition of our rule, where the notion of "confirmation" is explicated. For rule BAY_1, confirmation is increase in firmness, and our measure of increase in firmness is l, since it satisfies our preferred adequacy conditions.

In this section I have articulated a Bayesian rule in the form of a sentence schema, BAY_1, as the HTI requires. In doing so, I have made explicit the many choices we must make so that our logic can be articulated as a sentence schema with its corresponding side condition. This process allows us to situate several disagreements about Bayesianism in a common framework – disagreements about the notion of confirmation, the adequacy conditions on its measures, or our modelling choices, for example. As a consequence, we can better understand the commitments that are required in order to articulate a Bayesian rule and the impact of each commitment. Hopefully, the landscape of Bayesian inductive logics becomes a bit more clear when seen from the perspective of the HTI. In the next section I make explicit the ideal warranting conditions for the Bayesian rule I have articulated here, $\text{IWC}_{\text{BAY}_1}$.

4.4 Exposing the Ideal Warranting Conditions for our Bayesian Rule

In the previous section I articulated a Bayesian rule, BAY_1. In this section I expose the ideal warranting conditions for this rule, $\text{IWC}_{\text{BAY}_1}$. However, I will not be able to articulate $\text{IWC}_{\text{BAY}_1}$ in full detail. This is because the exact form of this schema will depend on our stance with respect to many open debates about models in science. These debates are beyond the scope of the current thesis. Therefore, I will sketch $\mathrm{IWC}_{\mathrm{BAY}_1}$ and point towards the philosophical work on models that must be done in order to fully develop this schema. By exposing $\mathrm{IWC}_{\mathrm{BAY}_1}$ we will realize that the accuracy of the information on inductive support that the instances of BAY_1 provide depends on the accuracy of the probability model being used.

Recall the central thesis of the HTI, in its general form, as introduced in chapter 2: given a rule R there is an associated sentence schema IWC_R such that, whenever $\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ can be formed,

$$v\,(\mathrm{R}\,(\mathrm{A}, \Phi_1, \ldots, \Phi_n)) = v\,(\mathrm{IWC}_\mathrm{R}\,(\mathrm{A}, \Phi_1, \ldots, \Phi_n))$$

Our rule of interest is BAY_1. In order to expose $\mathrm{IWC}_{\mathrm{BAY}_1}$, the question we need to answer is: what does the truthlikeness of the instances of BAY_1 depend on?

Let's try to answer this question. It is useful to begin by looking at one instance of BAY_1. I will now focus on argument [Die*].

$$\begin{array}{rl} P: & \text{A standard casino die has landed} \\ & \text{on an even number} \\ C: & \text{The die has landed on 2} \end{array} \qquad [\text{Die}^*]$$

Let us choose our previously developed model of a fair six-sided die, $\mathcal{M}_{1,[\mathrm{Die}]}$, as the probability model of our system. Thus, $\mathcal{M}_{1,[\mathrm{Die}^*]} = \mathcal{M}_{1,[\mathrm{Die}]}$. Standard casino die have certain features, like sharp edges and painted pips (not engraved), which ensure that they are as close to being fair as possible. However, these die are still imperfect. For instance, not all edges and corners are equally sharp, and the mass is not perfectly distributed inside the die. As a result, standard casino die may be very close to being fair, but they are not fair.

Thus, we can understand $\mathcal{M}_{1,[\mathrm{Die}^*]}$ as an idealized model of our system. Each outcome w of our standard casino die will actually have a probability other than 1/6, but it will be very close to it. By using $\mathcal{M}_{1,[\mathrm{Die}^*]}$ as our model, it is easier to deal with our system of interest. The advantages of using idealisations in our models may not be too evident here, given the simplicity of our system, but as our systems become more complex idealisations become increasingly necessary.

Idealisations make our model inaccurate and, therefore, $\mathcal{M}_{1,[\mathrm{Die}^*]}$ is an inaccurate model of our system. For now, all we need to acknowledge is that a model is inaccurate if it is not a perfect representation of its target system. I will say more about the idea of model accuracy shortly. Hence, $\mathcal{M}_{1,[\mathrm{Die}^*]}$ is an inaccurate model of our system, but we can come up with models of our standard casino die that are even more inaccurate. Take the following model $\mathcal{M}_{2,[\mathrm{Die}^*]}$, with $\Omega_{2,[\mathrm{Die}^*]} = \{1, 2, 3, 4, 5, 6\}$, $\mathcal{F}_{2,[\mathrm{Die}^*]} = 2^{\Omega_{2,[\mathrm{Die}^*]}}$, and

$$p_{2,[\text{Die}^*]}(w) = \begin{cases} 1/21 \text{ for } w = 1 \\ 2/21 \text{ for } w = 2 \\ 3/21 \text{ for } w = 3 \\ 4/21 \text{ for } w = 4 \\ 5/21 \text{ for } w = 5 \\ 6/21 \text{ for } w = 6 \end{cases}$$

$$P_{2,[\text{Die}^*]}(S) = \sum_{w \in S} p_{2,[\text{Die}^*]}(w)$$

In $\mathcal{M}_{2,[\text{Die}^*]}$, the probability mass function assigns probabilities to each outcome proportionally to the number of pips of that outcome. This probability mass function corresponds to a very unfair die. Thus, $\mathcal{M}_{2,[\text{Die}^*]}$ would be a very inaccurate model of a standard casino die.

Let's see, now, what rule BAY_1 says about argument [Die*] for each choice of model. Rule BAY_1 has three placeholders: E, H and $\mathcal{M}$. Placeholder E corresponds to the premise of argument [Die*], so I call it $E_{[\text{Die}^*]}$. Placeholder H corresponds to the conclusion of argument [Die*], so I call it $H_{[\text{Die}^*]}$. In both models $\mathcal{M}_{1,[\text{Die}^*]}$ and $\mathcal{M}_{2,[\text{Die}^*]}$, $E_{[\text{Die}^*]}$ corresponds to event $\{2,4,6\}$, $H_{[\text{Die}^*]}$ corresponds to event $\{2\}$ and $\neg H_{[\text{Die}^*]}$ corresponds to event $\{1,3,4,5,6\}$. Then,

$$l\left(H_{[\text{Die}^*]}, E_{[\text{Die}^*]}, \mathcal{M}_{1,[\text{Die}^*]}\right) \approx 0.4$$
$$l\left(H_{[\text{Die}^*]}, E_{[\text{Die}^*]}, \mathcal{M}_{2,[\text{Die}^*]}\right) \approx 0.28$$

Therefore:

> $\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{1,[\text{Die}^*]}\right)$: $E_{[\text{Die}^*]}$ confirms $H_{[\text{Die}^*]}$ to degree 0.4
>
> $\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{2,[\text{Die}^*]}\right)$: $E_{[\text{Die}^*]}$ confirms $H_{[\text{Die}^*]}$ to degree 0.28

Since a standard casino die is much closer to being a fair die than to being the die described by model $\mathcal{M}_{2,[\text{Die}^*]}$, $\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{1,[\text{Die}^*]}\right)$ is more truthlike than $\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{2,[\text{Die}^*]}\right)$. The truthlikeness of the instances of BAY_1 is determined by the accuracy of the probability model we use. Therefore:

> $\text{IWC}_{\text{BAY}_1}$: $\mathcal{M}$ is a perfectly accurate model of the system in argument A

The side condition for this schema must explicate the term "accurate" and it must also tell us how to fill in placeholders $\mathcal{M}$ and A. Placeholder $\mathcal{M}$ must be filled in with the probability model that we have chosen to represent our system of interest, and argument A must be filled in with the identifier of our argument of interest. Thus:

> $\text{IWC}_{\text{BAY}_1}\left([\text{Die}^*], \mathcal{M}_{1,[\text{Die}^*]}\right)$: $\mathcal{M}_{1,[\text{Die}^*]}$ is a perfectly accurate model of the system in argument [Die*]
>
> $\text{IWC}_{\text{BAY}_1}\left([\text{Die}^*], \mathcal{M}_{2,[\text{Die}^*]}\right)$: $\mathcal{M}_{2,[\text{Die}^*]}$ is a perfectly accurate model of the system in argument [Die*]

Since $\mathcal{M}_{1,[\text{Die}^*]}$ is a more accurate model of a standard casino die than $\mathcal{M}_{2,[\text{Die}^*]}$,

$$v\left(\text{IWC}_{\text{BAY}_1}\left([\text{Die}^*], \mathcal{M}_{1,[\text{Die}^*]}\right)\right) > v\left(\text{IWC}_{\text{BAY}_1}\left([\text{Die}^*], \mathcal{M}_{2,[\text{Die}^*]}\right)\right)$$

and, therefore, according to the HTI,

$$v\left(\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{1,[\text{Die}^*]}\right)\right) > v\left(\text{BAY}_1\left([\text{Die}^*], \mathcal{M}_{2,[\text{Die}^*]}\right)\right)$$

meaning that the degree of confirmation (as defined in the side condition of BAY_1) in argument $[\text{Die}^*]$ is closer to 0.4 than to 0.28. This result should be quite intuitive, and the HTI provides a clear logical picture of why this is so.

However, recall that the side condition of $\text{IWC}_{\text{BAY}_1}$ must also explicate the term "accurate". This poses an important challenge. What does it mean for a model to be accurate, and how do we assess its accuracy? Frigg and Nguyen (2017) call this challenge "the problem of standards of accuracy". These authors note that accuracy comes in degrees, it may be context dependent and it might make reference to the purposes of the model and the model user.

Addressing the problem of standards of accuracy quickly raises other questions about models. What are models? Maybe they are fictional objects (Godfrey-Smith, 2006), or set-theoretic structures (da Costa & French, 2000; Suppes, 1960), or descriptions (Levy, 2012, 2015). Our stance on the ontology of models may influence how we address the problem of standards of accuracy and many other debates (e.g. Frigg, 2010). Furthermore; what is it about a model that allows it to represent its target system? Maybe all it takes for a model $\mathcal{M}$ to represent a target system T is for a model user to stipulate that this is so (Callender & Cohen, 2006); maybe $\mathcal{M}$ must be similar to T in some respect (Giere, 2004); maybe $\mathcal{M}$ must be understood as a structure, and it must be isomorphic to the structure of T in order to represent it (Ubbink, 1960); or maybe $\mathcal{M}$ must allow competent users to draw some specific inferences about T (Hughes, 1997; Suárez, 2004). Our understanding of the representational relation between $\mathcal{M}$ and T will obviously influence how we define standards of accuracy for $\mathcal{M}$. For instance, if $\mathcal{M}$ represents T in virtue of being similar to it in some relevant respect, then the accuracy of this representation will be determined by the degree of similarity between $\mathcal{M}$ and T in the relevant respect. As Frigg and Nguyen (2017) note, model accuracy may also make reference to the purposes of the model and the model user, like Parker (2020) argues. Thus, the problem of standards of accuracy is complex and multifaceted, and addressing it is beyond the scope of this thesis.

Since I am not addressing the problem of standards of accuracy, I cannot provide a complete side condition for $\text{IWC}_{\text{BAY}_1}$, since I will not argue for any particular account of model accuracy. This means that I cannot provide a set of detailed examples to illustrate why $\text{IWC}_{\text{BAY}_1}$ is this particular schema, as I have done in the previous chapters. In particular, I cannot develop any measure of verisimilitude for the instances of $\text{IWC}_{\text{BAY}_1}$, since that would require a measure of model accuracy, which we are lacking. Therefore, I cannot show that $v(\text{BAY}_1(\text{A}, \mathcal{M})) = v\left(\text{IWC}_{\text{BAY}_1}(\text{A}, \mathcal{M})\right)$ independently of our interests. I hope, however, that $\text{IWC}_{\text{BAY}_1}$ is intuitive enough: $\text{IWC}_{\text{BAY}_1}$ captures the idea that, once we have fixed our understanding of confirmation in probabilistic terms, any given statement about the degree of confirmation in an argument will be more truthlike the more accurate our probability model is.

In this section I have exposed the ideal warranting conditions for rule BAY_1, IWC_{BAY_1}. In order to assess the verisimilitude of the instances of IWC_{BAY_1} we must be able to assess the accuracy of our probability models – that is, we face the "problem of standards of accuracy". The problem of standards of accuracy is complex and multifaceted. More work on model accuracy is required, and the HTI shows why and how this work is relevant in our understanding of BAY_1.

4.5 Ideal Warranting Conditions for All Bayesian Inductive Logics

In the previous section I exposed IWC_{BAY_1}; in this section I argue that IWC_{BAY_1} functions as IWC for *any* Bayesian inductive logic, not only for BAY_1. I will first provide a general argument for this point, and I will then quickly develop a Bayesian inductive logic other than BAY_1 to illustrate that IWC_{BAY_1} still functions as IWC for this logic. Showing that IWC_{BAY_1} functions as IWC for any Bayesian inductive logic is important in order to generalize the conclusion from the previous section, allowing us to show that the accuracy of our models plays an important role in our understanding of Bayesian inductive logics in general.

Any Bayesian inductive logic must articulate the notion of confirmation in probabilistic terms. Regardless of whether confirmation is understood as firmness or as increase in firmness, as a qualitative, comparative or quantitative notion, all Bayesian inductive logics must explicate confirmation in terms of probabilities. All these probabilities are given by, or calculated from, a probability model of our system of interest, $\mathcal{M}$. Therefore, the accuracy of any assessment of confirmation in a Bayesian inductive logic will depend on the accuracy of the probability model $\mathcal{M}$ being used.

Let me now illustrate the main point of this section with an example. Take the following Bayesian rule:

BAY_2: E confirms H to degree $P_{\mathcal{M}}(H|E)$

with the side condition that

(1) Confirmation is firmness;

(2) E is the premise of our argument of interest;

(3) H is the conclusion of our argument of interest;

(4) $P_{\mathcal{M}}$ is the probability function P in model $\mathcal{M}$;

(5) $\mathcal{M}$ is a probability model of our system, consisting of three elements, $(\Omega, \mathcal{F}, P)$, where Ω is the sample space, $\mathcal{F}$ the event space, and P a probability measure, $P : \mathcal{F} \to [0,1]$. These elements must satisfy Kolmogorov's axioms.

In order to expose IWC_{BAY_2} we have to ask: what does the truthlikeness of the instances of BAY_2 depend on? As expected, the truthlikeness of any instance of BAY_2 will depend on the accuracy of model $\mathcal{M}$. Let me illustrate this. I will focus again on argument [Die] and pick $\mathcal{M}_{1,[\text{Die}]}$ as the model of our system. Then,

$BAY_2\left([\text{Die}], \mathcal{M}_{1,[\text{Die}]}\right)$: $E_{[\text{Die}]}$ confirms $H_{[\text{Die}]}$ to degree $P_{1,[\text{Die}]}(H_{[\text{Die}]}|E_{[\text{Die}]})$

Given that we understand confirmation as firmness, if $\mathcal{M}_{1,[\text{Die}]}$ is a perfectly accurate model of our system then it is the case that "$E_{[\text{Die}]}$ confirms $H_{[\text{Die}]}$ to degree $P_{1,[\text{Die}]}(H_{[\text{Die}]}|E_{[\text{Die}]})$"; as our model becomes less accurate, this statement becomes less truthlike. Thus;

$$\text{IWC}_{\text{BAY}_2} = \text{IWC}_{\text{BAY}_1} : \mathcal{M} \text{ is a perfectly accurate model of the system in argument } A$$

In conclusion, since any Bayesian rule must make sense of confirmation in probabilistic terms, the truthlikeness of any instance of a Bayesian rule is determined by the accuracy of the probability model we use. Thus, for any Bayesian rule BAY_i, $\text{IWC}_{\text{BAY}_i} = \text{IWC}_{\text{BAY}_1} = \text{IWC}_{\text{BAY}}$, with IWC_{BAY} being the generic name I will use from now on for the Ideal Warranting Conditions of any Bayesian rule.

We can find precursors to this position in some classic literature. In his 1955 paper "The Problem of Inductive Inference", Neyman defends some of his ideas from Carnap's criticisms. Neyman's diagnosis is that most disagreements between Carnap and himself are rooted on a difference in emphasis: while Neyman rightly emphasizes the distinction between the phenomena of interest (what I have called "target systems") and the models of those phenomena, Carnap doesn't (Neyman, 1955, p. 14). This distinction allows Neyman to reach the following conclusion:

> [W]hatever the choice of the phenomena, the conclusions of a theory of inductive inference will always be applicable within the mathematical models of these phenomena and not within the domains of the phenomena themselves. Since, in many instances, the phenomena rather than their models are the subject of scientific interest, the transfer to the phenomena of an inductive inference reached within the model must be something like this: granting that the model *M* of phenomena *P* is adequate (or valid, or satisfactory, etc.) the conclusion reached within *M* applies to *P*. (Neyman, 1955, p. 17)

In this fragment, Neyman is exposing the ideal warranting conditions for probabilistic theories of induction, i.e., that a probability "model *M* of phenomena *P* is adequate (or valid, or satisfactory, etc.)".[17] This coincides with IWC_{BAY}. Thus, by failing to distinguish probability models from their target systems, Carnap fails to recognize the need for such ideal warranting conditions; this is another way to understand Neyman's diagnosis.

4.6 Implications

In this section I summarize some implications of thinking about Bayesian inductive logics from the perspective of the HTI. First, and crucially, we have seen that the HTI makes explicit the role of model accuracy in our understanding of inductive support within Bayesian inductive logics. By exposing IWC_{BAY} we realize that the accuracy of the information provided by any Bayesian rule depends on the accuracy of the probability model being used. More work is required in this area, and the HTI provides a clear framework for Bayesian logicians to draw on the rich literature on models and model accuracy.

[17] For our purposes, the differences between frequentism and Bayesianism do not matter, just like they do not matter for Neyman's purposes in this fragment either (Neyman, 1955, pp. 16–17).

Furthermore, the understanding of Bayesian inductive logics introduced so far has some epistemic consequences too. This is so because Bayesian inductive logics, like any inductive logic, describe the relation of support between two propositions, which has an impact on the relation of support between our beliefs in those propositions (de Grefte, 2020). For instance, *ceteris paribus*, the more a piece of evidence E supports a hypothesis H, the higher our credence in H should be. As introduced in section 4.2, the link between logic and epistemology is complex, and I say a bit more about it in the conclusion of this thesis. For now, remembering that there is a relation between these logical and epistemic matters is all that is required. As introduced in the previous chapter, this observation together with the HTI means that the ideal warranting conditions of a rule function both as a methodological and as a conceptual guide for researchers. Let's explore what this means for Bayesian rules.

The ideal warranting conditions of a rule function as a methodological guide to better inferences because they tell us which conditions we should maximize in order to maximize the truthlikeness of the information on inductive support that our rule provides. In the previous chapter we examined a rule which posited an entailment relation (rule CIS). For rules positing an entailment relation, maximizing the truthlikeness of their instances means maximizing the degree of support in our argument of interest, thus, maximizing the strength of our inferences.[18] Bayesian rules, however, do not posit an entailment relation but a weaker confirmation relation. For such rules, maximizing the truthlikeness of their instances does not guarantee stronger inferences; what it guarantees is access to more accurate information on inductive support, which allows for better inferences. This observation is undeniably vague – in order to make it more precise, more work on bridge principles for Bayesian inductive logics would be required, exploring the relation between confirmation and our epistemic attitudes. This is a matter for future work. For now, I hope the next paragraph illustrates what I mean in intuitive terms.

IWC_{BAY} functions as a methodological guide to better inferences in the following sense: in order to maximize $v(\text{BAY}_i(\text{A}, \mathcal{M}))$ we have to maximize $v(\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M}))$, and this is achieved by maximizing the accuracy of model $\mathcal{M}$. So, IWC_{BAY} tells us that we should maximize the accuracy of our models, since more accurate models will yield more truthlike information on inductive support, which will better inform our inferences. This is because the more accurate the information on inductive support that we have, the less guesswork is involved in correcting for deviations due to the model's inaccuracies. $\mathcal{M}_{1,[\text{Die}^*]}$ is a more accurate model of a standard casino die than $\mathcal{M}_{2,[\text{Die}^*]}$, hence, $l = 0.4$ is a more accurate estimate of the degree of confirmation in argument [Die*] than $l = 0.28$ and, therefore, it is a better guide for our inferences. This should sound quite intuitive. Of course, there are practical limitations to our modeling practices, meaning that most often we will use idealized models that fall short from perfect accuracy. This is where the next observation becomes relevant.

IWC_{BAY} also functions as a conceptual guide to assess the strength of our inferences: in order to assess $v(\text{BAY}_i(\text{A}, \mathcal{M}))$ we have to assess $v(\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M}))$, and this is achieved by assessing the accuracy of model $\mathcal{M}$. Hence, those factors which have an impact on the accuracy of our model also have an impact on $v(\text{BAY}(\text{A}, \mathcal{M}))$, i.e., on the truthlikeness of the information on inductive support provided by any Bayesian rule using model $\mathcal{M}$. Knowing that the information on inductive support we obtain using model $\mathcal{M}$ is more (or less) truthlike is relevant in adjusting the

[18] Since, *ceteris paribus*, stronger arguments result in stronger inferences.

strength of our inferences using model $\mathcal{M}$.[19] This should sound quite intuitive too. For instance:

> If you build, for example, a statistical model of financial markets where you neglect the effects of psychological chain reactions on asset prices (e.g., because the model would get too complicated), you will underestimate the probability of a stock exchange crash. Our actual degrees of belief should, to some extent, compensate for such modeling decisions. (Sprenger and Hartmann, 2019, p. 322)

In this scenario, our model of financial markets is inaccurate – it neglects the effects of psychological chain reactions on asset prices – and this inaccuracy results in an inaccurate (underestimated) probability of a stock exchange crash. As Sprenger and Hartmann suggest, then, our degrees of belief should compensate for such modeling decisions.

These epistemic consequences of the HTI are in line with some recent suggestions in the literature on Bayesian epistemology. Morey et al. (2013) and Gelman and Shalizi (2013), for example, note that most philosophical literature on Bayesian epistemology does not take into account the importance of checking one's own models. However, "most practising Bayesians worry about the appropriateness of their models and hence engage in model checking" (Morey et al., 2013, p. 69). In fact, Gelman and Shalizi note the following: "[s]ince we are quite sure our models are wrong, we need to check whether the misspecification is so bad that inferences regarding the scientific parameters are in trouble" (Gelman and Shalizi, 2013, p. 17).[20] And, as developed above, the need for model-checking follows from looking at Bayesianism from the perspective of the HTI; IWC_{BAY} tells us that those factors which have an impact on the accuracy of our model are relevant in assessing the strength of our inferences. Thus, we must check our models in order to assess their accuracy, so we can assess the strength of our inferences. The HTI offers a logical underpinning for a Bayesian epistemology that includes model-checking – the kind of Bayesian epistemology that Morey et al. call for.

It might be illustrative to see how Morey et al. (2013) suggest that we go about model-checking to assess the accuracy of our models. Theirs is, of course, just one of many possible approaches. Morey et al. start by suggesting that the representational relation between a model and its target system can be understood using Hughes (1997) "denotation, demonstration, interpretation" (DDI) framework, sketched in figure 4.1 (from (Morey et al., 2013, p. 71)). As introduced in §4.4, Hughes' is an inferential account of the representational relation between models and their target systems.

Morey et al. then focus on the final step of the process shown in the DDI model, that of "interpretation". This is the step in which the statistical inferences can be interpreted into real-world inferences, and this is where model-checking comes into play. Morey et al. then acknowledge the usefulness of the model-checking methods suggested by Gelman and Shalizi (2013) (despite disagreeing with their understanding of the nature of the model-checking process). Here is a brief summary of the model-checking process following Gelman and Shalizi's proposed methods: "the hypothesized model makes certain probabilistic assumptions, from which other probabilistic implications follow deductively. Simulation works out what those implications are, and tests check

[19] In general, an inaccurate model can result both in an underestimation or an overestimation of the degree of confirmation in our argument of interest.

[20] Although, as Morey et al. point out, models, as representations, cannot be *right* or *wrong*. This is why I talk in terms of model *accuracy*, following Frigg and Nguyen (2017).

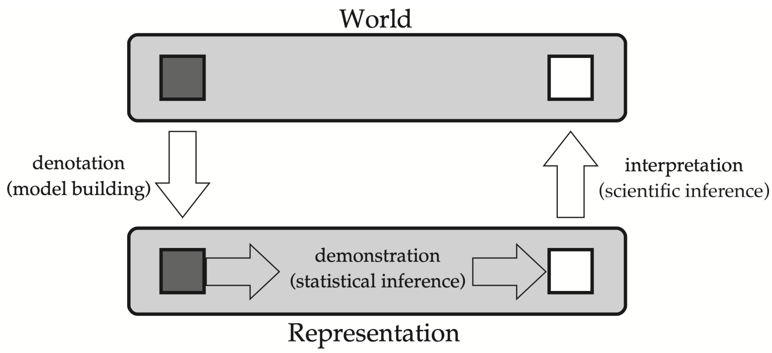

Figure 4.1: "Hughe's (1997) DDI model of scientific representation. 'World' squares represent phenomena (dark) or propositions about phenomena (light); 'representation' squares represent models (dark) or inferences about models (light)." (Morey, Romeijn, and Rouder, 2013, p. 71)

whether the data conform to them" (Gelman and Shalizi, 2013, p. 18). Thus, according to these authors, it is by checking the data against the model's implications that we can assess the accuracy of the model and, then, interpret the statistical inferences into real-world inferences. Of course, much more must be said about this process, but hopefully this sketch suffices to visualize how one might tackle the problem of standards of accuracy.

The epistemic consequences of viewing Bayesianism from the perspective of the HTI are also in line with some recent suggestions by Sprenger and Hartmann (2019, pp. 311–26). In their "Variation 12: Models, Idealizations and Objective Chance", Sprenger and Hartmann motivate and develop a suppositional analysis of probability in Bayesian inference. According to this view, $P(H)$ is not an unconditional but a conditional model-relative probability. In fact, the authors argue, we should always write $P_{\mathcal{M}}(H)$ since "all probabilities in Bayesian inference are relative to a statistical model" (Sprenger and Hartmann, 2019, p. 320). This suppositional analysis of probability has some interesting consequences:

> How does the suppositional analysis link model predictions to the real world? This question is especially urgent in contexts such as climate science, where doubts about the adequacy of large-scale models for making probabilistic predictions are widespread. [...] In general, confidence in the predictions of a statistical model $\mathcal{M}$ depends on whether its constituents capture the relevant aspects of the target system. When the target system is highly complex and hard to predict, we will have to preserve a healthy dose of skepticism toward the predictions of $\mathcal{M}$. On the other hand, the better the constituents of $\mathcal{M}$ describe relevant aspects of the target system, the more justification do we have for an inference by analogy, and for transferring their predictions to our actual epistemic attitudes. In this respect, Bayesian inference is just another form of model-based reasoning in science: its soundness depends on whether the overall model $\mathcal{M}$ is well chosen or inadequate. (Sprenger and Hartmann, 2019, pp. 321–22)

Thus, the epistemic consequences of viewing Bayesian inductive logics from the perspective of the HTI coincide with Sprenger and Hartmann's suggestions. In this way, the HTI also offers a logical underpinning for Sprenger and Hartmann's conception of Bayesian inference as a form of model-based reasoning. There are further connections to explore between the account of Bayesian inductive logics introduced here and Sprenger and Hartmann's views, especially regarding the co-ordination principles bridging logical and epistemic attributes. This is a matter for future work.

The HTI is also useful in developing better bridge principles for inductive logic. For instance,the analysis of Bayesian inductive logics introduced here allows us to generalize Fitelson's suggested bridge principle for these logics (Fitelson, 2006). Fitelson is interested in "the relation between inductive logic (the confirmation as increase in firmness relation) and epistemology (the relation of incremental evidential support)" (Fitelson, 2006, p. 500). Fitelson does not explicitly introduce the notion of "incremental evidential support", but this is not relevant for the point I wish to make; it will suffice to remember that it is an epistemic concept. Fitelson then suggests the following bridge principle: if an agent a knows that E, and a knows that $l(H, E, \mathcal{M}) > 0$, and a knows that model $\mathcal{M}$ is a correct model of the stochastic process S that generated the event E, then a knows that E evidentially supports H (Fitelson, 2006, pp. 508–509).[21] The core idea behind this bridge principle is that the information on confirmation that we obtain by using model $\mathcal{M}$ only bears on our epistemic attitudes if we know model $\mathcal{M}$ to be correct (or, in my terms, "perfectly accurate"). This suggestion is clearly on the right path, but requiring S to know that model $\mathcal{M}$ is correct seems too strong. In terms of the HTI, this is equivalent to requiring S to know that $\mathrm{IWC_{BAY}}(\mathrm{A}, \mathcal{M})$ is true . However, as I have argued in this section, this is not necessary. If this were necessary, then the fact that $l(H_{[\mathrm{Die}^*]}, E_{[\mathrm{Die}^*]}, \mathcal{M}_{1,[\mathrm{Die}^*]}) \approx 0.4$ would not bear on our epistemic attitudes, since we actually know model $\mathcal{M}_{1,[\mathrm{Die}^*]}$ to not be correct (or perfectly accurate) because it is an idealized model. But we know model $\mathcal{M}_{1,[\mathrm{Die}^*]}$ to be very accurate and, in particular, to be much more accurate than $\mathcal{M}_{2,[\mathrm{Die}^*]}$; this constrains our epistemic attitudes. Thus, for a given piece of information on confirmation to bear on our epistemic attitudes, agent a does not need to know that $\mathcal{M}$ is the "correct" model of our target system. Instead, we can weaken this bridge principle by requiring a to assess the accuracy of model $\mathcal{M}$, which is done by model-checking (Gelman & Shalizi, 2013; Morey et al., 2013). The more accurate the model, the more truthlike the information on confirmation it provides, and this is epistemically useful. In conclusion, the HTI helps us realise that the information on confirmation that a Bayesian rule provides is still epistemically relevant despite our model not being perfect, and to develop bridge principles that take this into account.

Some readers may still be inclined to think that a Bayesian inductive logic can simply *assume* a probability model, and work from there. According to this criticism, then, model accuracy is irrelevant for Bayesian inductive logics, and it is, if anything, an epistemic matter. This is not so. Bayesian inductive logics, like all inductive logics, are accounts of the relation of support between propositions in an ampliative argument. For instance, if we are interested in argument [Die*], a Bayesian inductive logic must tell us how much the proposition that "a standard casino die has landed on an even number" supports the proposition that "the die has landed on 2". These are propositions about the world. Thus, in order to understand the relation of support between these propositions it does not suffice to make the right operations on a given probability model.

[21] Again, I do not think we can talk of correct or incorrect models, since models are representations and representations are not correct or incorrect, but we can substitute Fitelson's term "correct" for my term "perfectly accurate".

The model is not itself our system of interest.[22] We need the tools to translate our information about that model to information about the world: the HTI makes this need salient and gives it its rightful place in our logical framework.

Of course, much work in Bayesian inductive logic can (and must) proceed without worrying about model accuracy. We might take the objective probabilities as given, and then we might still ask: how do we understand and measure confirmation? This kind of work is what we might call "first-order" work. It is the work required to develop a rule of induction, which provides information on inductive support. The articulation of rule BAY_1 in this chapter is an example of how one might go about this, and the kind of decisions involved. However, the HTI shows that we need to do some second-order work too. That is, we also need to ask: how *accurate* is this information about inductive support? The HTI gives us the beginning of an answer: it depends on the accuracy of our probability model.

In this section I have summarized some of the implications of understanding Bayesian inductive logics from the perspective of the HTI. Crucially, the HTI makes explicit the role of model accuracy in our understanding of inductive support within Bayesian inductive logics: by exposing IWC_{BAY} we realize that the accuracy of the information provided by any Bayesian rule depends on the accuracy of the probability model being used. This perspective on Bayesian inductive logics has some epistemic consequences too. As introduced in the previous chapter, a rule's ideal warranting conditions can function both as a methodological and as a conceptual guide for researchers. IWC_{BAY} functions as a methodological guide in the following sense: it tells us that we should maximize the accuracy of our models, since more accurate models will yield more truthlike information on inductive support, which will better inform our inferences. This is because the more accurate the information on inductive support that we have, the less guesswork is involved in correcting for deviations due to the model's inaccuracies. IWC_{BAY} also functions as a conceptual guide to the factors that are relevant in assessing the strength of our inferences: it tells us that those factors which have an impact on the accuracy of our model are relevant in assessing the strength of our inferences. These epistemic consequences are in line with some recent suggestions in the literature on Bayesian epistemology, like (Gelman & Shalizi, 2013), (Morey et al., 2013) and (Sprenger & Hartmann, 2019, pp. 311–26). In this way, the HTI offers a logical underpinning for these epistemic positions. Furthermore, the HTI is also useful in developing better bridge principles for inductive logics; the analysis of Bayesian inductive logics introduced in this chapter allows us to generalize Fitelson's bridge principle for these logics (Fitelson, 2006), helping us realize that the information on confirmation that we obtain by using a given model is still epistemically relevant despite our model not being perfect.

4.7 Conclusion

In this chapter I have explored Bayesian inductive logics from the perspective of the HTI. I have begun by expressing a specific Bayesian inductive logic, BAY_1, as a sentence schema with its corresponding side condition. In doing so, I have situated several debates about Bayesian inductive logics in a common framework; debates about the notion of confirmation, the adequacy conditions

[22]See (Morey et al., 2013, p. 74) or (Neyman, 1955, p. 17) for similar observations.

for measures of confirmation or our modelling choices. This clarifies the impact of these debates on our Bayesian apparatus and the connections between the issues they address. Hopefully, the landscape of Bayesian inductive logics becomes a bit more clear when seen from the perspective of the HTI. I have then exposed the ideal warranting conditions for any Bayesian inductive logic, IWC_{BAY}. This schema captures the idea that the accuracy of the information on inductive support that any Bayesian inductive logic provides is determined by the accuracy of the probability model being used.

Finally, I have explored some epistemic implications of understanding Bayesian inductive logics from the perspective of the HTI. IWC_{BAY} functions as a methodological guide to better inferences; it tells us that we should maximize the accuracy of our models, since more accurate models will yield more truthlike information on inductive support, which will better inform our inferences. IWC_{BAY} also functions as a conceptual guide to assess the strength of our inferences; it tells us that the factors which are relevant for the accuracy of our probability model are relevant in assessing the strength of our inferences. What exactly those factors are will depend on how we resolve the problem of standards of accuracy. These epistemic consequences are in line with Morey's call for a Bayesian epistemology that accounts for the practice of model-checking, and with Sprenger and Hartmann's understanding of Bayesian inference as a form of model-based reasoning. Thus, the HTI provides a logical underpinning for these epistemic positions. Furthermore, understanding Bayesian inductive logics from the perspective of the HTI can help us develop better bridge principles for these logics. I have illustrated this by generalizing Fitelson's (2006) bridge principle for a Bayesian inductive logic.

Conclusion

In this thesis I have motivated and developed a Hybrid Theory of Induction (HTI) and I have explored some of its virtues and implications. The HTI is a hybrid second-order model of inductive support. It is a *hybrid* model of inductive support because it holds that two ingredients play a necessary role in understanding inductive support: rules and facts. It is a *second-order* model of inductive support because it is a model within which first-order models of inductive support (i.e. logics of induction) can fit. In chapter 1 I argued that we need both rules and facts to play a role in a successful account of inductive support. Rules of induction accurately describe relations of inductive support when they are warranted; facts do the warranting work. I called this type of warrant "factual warrant". The resulting account is both functional and accurate, it helps us make sense of how different rules of induction can coexist and it allows us to resolve some current debates on induction (like the apparent divide between Norton and rule-based theorists). For the purposes of chapter 1 I adopted the MTI as an account of factual warrant. In chapter 2 I developed a Graded account of Factual Warrant (GFW), according to which factual warrant comes in degrees. I integrated the GFW in the HTI. I then showed that the GFW has some advantages over the MTI: it illuminates the connection between factual warrant and inductive support, and it can successfully account for the role of idealisations and theory in our understanding of inductive support. In chapter 3 I argued that the HTI is also useful for agents, since it can provide methodological guidance to ensure strong inferences and conceptual guidance to assess the strength of our inferences. Finally, in chapter 4, I explored Bayesian inductive logics from the perspective of the HTI. This analysis brought to light the central role that probability models play in Bayesian inductive logics, offering a logical underpinning for some recent suggestions in Bayesian epistemology. Furthermore, throughout this thesis I have analysed in detail three rules of induction from the perspective of the HTI: enumerative induction in chapter 2, causal inference in chapter 3 and Bayesian inductive logics in chapter 4. These analyses have illustrated how the HTI can help us think more clearly about rules of induction, offering new tools to tackle existing challenges.

Many questions, however, remain open. A central one has to do with the nature of inductive support. I have been talking about inductive support for this whole thesis, but what *is* it? This is, of course, an important question. Still, as I anticipated in the introduction of this thesis, I have not discussed the metaphysics of inductive support because it is not required for my project. In this thesis I have been concerned with modelling inductive support. Just like we can advance our models of causation without having a clear metaphysics of causation, so we can advance our models of inductive support without having a clear metaphysics of inductive support. We want our model of inductive support to behave in certain ways, independently of what the nature of inductive support is. We want our model, for example, to tell us that the degree of support in argument

[Bismuth] is much greater than in argument [Wax], and to illuminate why this is so. Throughout this thesis I have exposed some of the virtues of the HTI as a model of inductive support, and I have summarized them in the preceding paragraph. These virtues are all independent of how we choose to understand the nature of inductive support.

A different concern has to do with the coexistence of several rules of induction. What happens if different rules say different things about the degree of support in an inductive argument? This question is not unique to the HTI. In fact, the only way to escape this question is to defend the existence of some universal rule of induction that somehow dominates over all others, and the prospects for this are dim, as I and many others have argued. Thus, we must come to terms with the idea that different rules may say different things about the same argument. I believe the right way to make sense of this is to be pluralists about inductive support. As I argued in chapter 4, even within Bayesian rules we might favour different measures of inductive support depending on the adequacy conditions we impose on such measures. Thus, even once we are committed to understanding inductive support in terms of probabilities, there are different ways in which we can explicate inductive support, each of which captures a different relation between E and H. This lesson is extensive to non-Bayesian rules of induction. Each rule captures a particular mode of support, that is, a particular way in which a piece of evidence can support a hypothesis. More work is required to make sense of these possibly conflicting judgements of inductive support. This is the kind of work that would benefit from a clear metaphysics of inductive support. If inductive support is to be understood in terms of possible worlds, for example, different rules of induction may be carving out the space of possibilities differently. This is a matter for future work.

This thesis has made salient the need for further work in several other areas. More work on verisimilitude is required if we are to develop good measures of verisimilitude for the instances of the ideal warranting conditions of our rules of interest. I briefly explored some of the challenges related to developing such measures in chapter 3. A second important area requiring further work is the problem of standards of accuracy, introduced in chapter 4. As I discussed in that chapter, we need a better understanding of model accuracy if we are to develop good measures of verisimilitude for the ideal warranting conditions of our Bayesian inductive logics. In turn, we need such measures if we are to assess the verisimilitude of the information on inductive support that our Bayesian rules provide.

So far I have outlined some of the questions that the HTI raises and some areas in need of further work; let me now point to some future work in induction where the HTI can help.

Bridge Principles for Inductive Logics

The HTI can be useful in exploring bridge principles for inductive logics, as I already gestured to in chapter 4. Much literature on bridge principles draws on MacFarlane's influential talk "In What Sense (If Any) Is Logic Normative for Thought?", delivered at the 2004 Central Division APA symposium on the normativity of logic (MacFarlane, 2004). MacFarlane uses the term "bridge principle" to refer to any "general principle that articulates a substantive relation between 'facts' about logical consequence (or perhaps an agent's attitudes towards such facts) on the one hand, and norms governing the agent's doxastic attitudes vis-à-vis the propositions standing in these logical relations on the other" (Steinberger, 2022, §4). Since MacFarlane was interested in the relation between *deductive* logic and thought, his bridge principles are bridging facts about *logical consequence* with norms governing the agent's doxastic attitudes. MacFarlane offered a blueprint to generate such bridge principles:

$$\text{If } E \models H \text{ (or } \alpha(E \models H)) \text{ then } N(\beta(E), \gamma(H))$$

where α, β and γ are doxastic attitudes and N is a normative claim.[23] Thus, according to MacFarlane, bridge principles for deductive logic take the form of a conditional. The antecedent of this conditional is either a fact about logical consequence (that $E \models H$) or an agent's doxastic attitude α towards such fact (that agent *a believes* that $E \models H$, for example). The consequent of this conditional is a normative claim about the agent's doxastic attitudes towards E and H. Here is one bridge principle we can generate using this blueprint: if an agent believes that $E \models H$ then if the agent believes that E she ought to believe that H.

MacFarlane offers three parameters of this blueprint that can be varied to generate the space of possible bridge principles (MacFarlane, 2004, p. 6):

1. *Type of deontic operator*. Do facts about logical validity give rise to strict *obligations*, *permissions*, or (defeasible) *reasons* for belief? We can refer to any one of these deontic operators as O.

2. *Polarity*. Are these obligations/permissions/reasons *to believe*, or merely *not to disbelieve*?

3. *Scope of deontic operator*. The normative claim N is a conditional: what one ought/may/has reason to believe with respect to H depends somehow on what one believes, or ought/may/has reason to believe, with respect to E. Does the deontic operator O govern the consequent of the conditional, $\beta(E) \supset O(\gamma(H))$, or both the antecedent and the consequent, $O(\beta(E)) \supset O(\gamma(H))$, or the whole conditional, $O(\beta(E) \supset \gamma(H))$?

Different choices in these parameters, or in the doxastic attitudes α, β and γ, will result in different bridge principles (see (MacFarlane, 2004, p. 7) for a list of all possible combinations).[24]

MacFarlane's blueprint allowed for a systematic study of the space of bridge principles for deductive logic (see (Steinberger, 2022, §4) for some relevant literature), but we are lacking a systematic study of the space of bridge principles for inductive logics. The HTI can help us here, by allowing us to generalize MacFarlane's blueprint for the case of inductive logics. We must first generalize the notion of bridge principle. Following MacFarlane and Steinberger, we might say the following:

> A bridge principle is a general principle that articulates a substantive relation between "facts" about the logical relations of support between propositions (or perhaps an agent's attitudes towards such facts) on the one hand, and norms governing the agent's doxastic attitudes vis-à-vis the propositions standing in these logical relations on the other.

[23] Doxastic attitudes are, for example, *belief*, *disbelief*, or *degree of belief*.

[24] MacFarlane lists all possible combinations of bridge principles where the antecedent is not attitudinal (it is a fact about logical consequence, not an agent's doxastic attitude towards that fact) and the doxastic attitudes in the consequent are "belief". That is, he only varies the type of deontic operator, the polarity of the normative claim and the scope of the deontic operator. It is easy to generate other bridge principles for different doxastic attitudes.

This definition of bridge principle is more general than the one offered by Steinberger, which specifies that the first relata in a bridge principle is a fact about "logical consequence". In general, facts about the relations of support between propositions (whatever their strength) bear on our doxastic attitudes towards those propositions.

We can now rely on the HTI to articulate a blueprint for bridge principles for inductive logics, following MacFarlane's model. Just like the antecedent of MacFarlane's blueprint contains a "fact" about logical consequence, we want the antecedent of our blueprint to contain a "fact" about inductive support. However, in contrast with rules of deduction, rules of induction are not universal. The truthlikeness of the information on inductive support that a given rule of induction provides will vary between contexts. Hence, the information on inductive support that a given rule provides, $\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$, must be accompanied by a measure of its truthlikeness, $v(\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n))$ (or the agent's doxastic attitude towards $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$).[25] Hence, our blueprint will look like this:

$$\text{If } F(\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n), \mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)) \text{ then } N(\beta(E), \gamma(H))$$

where E is the premise of argument A, H is the conclusion of argument A, and F is some claim about $\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ and $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$, which might include doxastic attitudes towards these sentences. For instance, here is one bridge principle we might generate using this blueprint:[26]

> If EIS(A) says that E entails H, and S's credence in $\mathrm{IWC_{EIS}}(\mathrm{A})$ is k, then if S believes that E, S's credence in H should be k.

I am not endorsing this bridge principle, but only illustrating the kind of bridge principles that this blueprint allows us to generate.

A central lesson from the HTI is that some claim about $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ must play a role in our bridge principles for inductive logic. This is in line with the epistemic consequences of the HTI that I explored in chapters 3 and 4. It is also in line with the way many authors understand the relation between a rule of induction and the inferences we make with it. Let us look at some examples.

For instance, in the context of causal inference in comparative group studies, Fuller is committed to the idea that the balance in the total contribution of confounders, $v(\mathrm{IWC_{CIS}}(\mathrm{A}))$, impacts our confidence in the conclusion of argument A. Fuller says the following: "[t]he various techniques and tricks used in comparative studies – randomization, double blinding, stratification, matching – are an attempt to bring the distribution of [confounders] closer to the ideal, so that we can feel more confident that any difference in outcome is due to the exposure" (Fuller, 2019, p. 920). Fuller's ideal distribution of confounders is equivalent to $\mathrm{IWC_{CIS}}$. In this fragment, Fuller is expressing a connection between a fact about the balance in the contribution of confounders and a fact about our doxastic attitude (our credence) towards the conclusion of our argument of interest. In particular, Fuller is telling us that we should be more confident in the conclusion

[25] Terms $\mathrm{R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ and $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ are introduced in chapter 2. Someone with externalist inclinations will probably want $v(\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n))$ to play a role in our bridge principles; someone with internalist inclinations will probably prefer to assign the same role to the agent's credence in $\mathrm{IWC_R}(\mathrm{A}, \Phi_1, \ldots, \Phi_n)$ instead. Our blueprint allows for both options.

[26] Rule EIS is introduced in chapter 2.

of our causal inferences in comparative group studies as the total contribution of confounders is more balanced, that is, as $v(\text{IWC}_{\text{CIS}}(\text{A}))$ increases. Such a connection between $\text{IWC}_{\text{CIS}}(\text{A})$ and our confidence in the conclusion of argument A will be reflected in any bridge principle for rule CIS that we articulate using our blueprint, since $v(\text{IWC}_{\text{CIS}}(\text{A}))$ (or some doxastic attitude towards $\text{IWC}_{\text{CIS}}(\text{A})$) will always be in the antecedent of the bridge principle, and the agent's doxastic attitude towards the conclusion of argument A will be in the consequent of the bridge principle. For example, here is one bridge principle for rule CIS that we can articulate using our blueprint:[27]

> If CIS(A) says that E entails H, and $v(\text{IWC}_{\text{CIS}}(\text{A})) = k$ then if S believes that E, S's credence in H should increase with k.

This bridge principle has an externalist flavour, since its antecedent only contains facts about CIS(A) and $\text{IWC}_{\text{CIS}}(\text{A})$, but no doxastic attitudes. We can also generate a similar bridge principle with an internalist flavour:

> If S knows that CIS(A) says that E entails H, and S's credence in $\text{IWC}_{\text{CIS}}(\text{A})$ is k, then if S believes that E, S's credence in H should increase with k.

Similarly, as I discussed in chapter 4, in the context of Bayesian epistemology Sprenger and Hartmann are committed to the idea that the accuracy of our model $\mathcal{M}$, $v(\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M}))$, impacts our confidence in the conclusion of argument A. Sprenger and Hartmann tell us, for example, that "the better the constituents of $\mathcal{M}$ describe relevant aspects of the target system, the more justification do we have [...] for transferring their predictions to our actual epistemic attitudes" (Sprenger and Hartmann, 2019, p. 322). Such a connection will be reflected in any bridge principle for Bayesian inductive logics that we generate using the blueprint I have suggested, since $v(\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M}))$ (or some doxastic attitude towards $\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M})$) will always be in the antecedent of the bridge principle, and the agent's doxastic attitude towards the conclusion of argument A will be on the consequent of the bridge principle. Suppose, for instance, that we understand confirmation as firmness, so that we are interested in rule BAY_2.[28] Here is one bridge principle we can generate for this rule:

> If $\text{BAY}_2(\text{A}, \mathcal{M})$ says that E supports H to degree $P_{\mathcal{M}}(H|E)$, and $v(\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M})) = k$, then if S believes that E, S's credence in H should increase with k, with S's maximum credence in H being $P_{\mathcal{M}}(H|E)$ when $\text{IWC}_{\text{BAY}}(\text{A}, \mathcal{M})$ is true.

In conclusion, the HTI allows us to articulate a blueprint to generate bridge principles for inductive logics. Crucially, both $\text{R}(\text{A}, \Phi_1, \ldots, \Phi_n)$ and $\text{IWC}_{\text{R}}(\text{A}, \Phi_1, \ldots, \Phi_n)$ must play a role in the antecedent of such principles. Such a blueprint allows for a systematic study of the space of possible bridge principles for inductive logics. Exploring the space of possible bridge principles generated by this blueprint is a matter for future work.

Non-Epistemic Values and Inductive Support

The HTI might also help us explore the role of non-epistemic values in scientific reasoning. Non-epistemic values are widely considered to play a legitimate role in some "external" parts of sci-

[27] Rule CIS is introduced in chapter 3.

[28] Rule BAY_2 is introduced in chapter 4.

ence, like the choice of hypotheses to pursue (Longino, 1990). However, the question of whether non-epistemic values can play a legitimate role in scientific reasoning itself is more controversial. Heather Douglas is an influential voice in this debate. Douglas has convincingly argued that "non-epistemic values are required in science wherever non-epistemic consequences of error should be considered" (Douglas, 2000, p. 559). Thus, according to Douglas, non-epistemic values may also play a legitimate role in each of the three "internal" stages of science: choice of methodology, gathering and characterization of the data, and interpretation of the data. However, Douglas does not let non-epistemic values cut too deep into the workings of scientific reasoning: "Hempel was right in asserting that whether or not a piece of evidence is confirmatory of a hypothesis (given a set of background assumptions) is a relationship in which value judgments have no role" (Douglas, 2000, p. 565). Relations of support between evidence and hypotheses are shielded from non-epistemic values, say Hempel and Douglas. However, the HTI casts some doubt on this claim.

The HTI shows that the degree of support in an argument depends on our interests. For example, EIS([Bismuth]) says that *P* entails *C* in argument [Bismuth], so that the degree of support is maximal. If EIS([Bismuth]) is true, then it is the case that *P* entails *C* in argument [Bismuth]; as EIS([Bismuth]) becomes less truthlike, then it becomes less truthlike that the degree of support in argument [Bismuth] is maximal, which means that the actual degree of support in argument [Bismuth] decreases. But verisimilitude is dependent on our interests (Northcott, 2013, pp. 1481–82), hence, the degree of support in an argument depends on our interests.

Let me illustrate this point with a thought experiment I introduced in chapter 2. The HTI tells us that $v(\text{EIS}([\text{Bismuth}]) = v(\text{IWC}_{\text{EIS}}([\text{Bismuth}]))$. Hence, the actual degree of support in argument [Bismuth] depends on $v\,(\text{IWC}_{\text{EIS}}([\text{Bismuth}]))$. $\text{IWC}_{\text{EIS}}([\text{Bismuth}])$ is the sentence "all samples of bismuth are like the observed samples of bismuth with respect to their melting temperature". Imagine a scenario, call it scenario 1, where every unobserved sample of bismuth melts at a different temperature within 10^{-30}ºC of 271ºC, but none melts exactly at 271ºC. Imagine another scenario, call it scenario 2, where all unobserved samples of bismuth melt exactly at 271ºC except for one, which melts at 5,000ºC. In both scenarios, the observed samples of bismuth melt at exactly 271ºC. Is IWC_{EIS} ([Bismuth]) more truthlike in scenario 1 or in scenario 2? The answer to this question, of course, will depend on our interests. Do we only care about whether any given sample melts exactly at 271ºC? If a sample doesn't melt at 271ºC, do we care about its actual melting temperature? If so, do we care about big differences much more than about small differences? We must answer these questions in order to develop a measure of verisimilitude that is properly capturing our interests, goals and purposes (see §2.6.2 and §2.6.3 for some examples of such measures). Answering these questions requires non-epistemic value judgements. Hence, the HTI suggests that non-epistemic values may have a deeper influence in scientific reasoning than previously recognized, and it provides a useful framework to explore their role.

Further Applications

I hope the HTI can be useful in other areas of the history and philosophy of science. For instance, we can make use of the HTI to reconstruct cases of good scientific reasoning more clearly. Successful episodes of scientific reasoning are usually described from the perspective of a particular rule of induction – one that works well in that scenario. Considerations about the limitations of that rule, or the reasons why it is suitable in that scenario, are often ignored or left implicit. These considerations, however, are relevant to the scientists and to those who reconstruct their inferences. The HTI makes these considerations explicit. In the context of Bayesian inductive logics, for example, this perspective allows us to clarify the role that the accuracy of our probability models play in understanding inductive support, as well as the role of idealised models and the need for model-checking. This allows us to better reconstruct the kind of inferential practices that Bayesians are actually involved in (Gelman & Shalizi, 2013).

Relatedly, as I anticipated in chapter 1, the HTI may also allow us to respond to Lipton's "question of description", which asks what principles we use in making inductive inferences. After examining several influential rules of induction in the literature, Lipton concludes that all of them are "finding inductive support where there is none and overlooking cases of genuine support" (Lipton, 2004, pp. 17–18). None of those rules is, by itself, a good description of our inferential practices. According to Lipton, none of these approaches gives "enough structure to the black box of our inductive principles to determine the inferences and judgements we actually make" (Lipton, 2004, p. 18). The HTI can help us here, as I already suggested in chapter 1. I can now provide a slightly more developed answer to this question: we use rules of induction that we believe are strongly warranted in our context. We may choose to use rule EIS whenever we have good reasons to believe that $v(\mathrm{IWC}_{\mathrm{EIS}}(\mathrm{A}))$ is very high; that is, whenever we have good reasons to believe that all entities of interest are like those observed with respect to the property being generalized. If, instead, we have access to a probability model of our target system, and we have good reason to believe that it is an accurate model, so that $v(\mathrm{IWC}_{\mathrm{BAY}_i}(\mathrm{A}, \mathcal{M}))$ is high, we may choose to use a Bayesian rule to guide our inferences. In this way, we ensure that our inferences are as well-informed as possible. This is, of course, just a preliminary answer, but I believe the HTI provides the right tools to respond to Lipton's question of description.

* * *

These concluding remarks have highlighted some of the questions that remain open regarding the subject matter of this thesis, some areas in need of further work, and some potential applications of the HTI. It is my hope that the HTI will allow us to explore these and many other topics from a new perspective, serving as a framework to think about induction more productively.

Bibliography

Aliseda, A. (2006). *Abductive Reasoning. Logical Investigations into Discovery and Explanation*. Springer.

Baker, A. (2020). Schemas for Induction. *Studies in History and Philosophy of Science Part A*, *82*, 114–119.

Bird, A., & Tobin, E. (2022). Natural Kinds. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2022). Metaphysics Research Lab, Stanford University.

Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press.

Britton, A., McKee, M., Black, N., McPherson, K., Sanderson, C., & Bain, C. (1999). Threats to Applicability of Randomised Trials: Exclusions and Selective Participation. *Journal of Health Services Research & Policy*, *4*(2), 112–121.

Burger, I. C., & Heidema, J. (2005). For Better, for Worse: Comparative Orderings on States and Theories. In *Confirmation, Empirical Progress, and Truth Approximation* (pp. 459–488, Vol. 83). Rodopi.

Byar, D. P., Simon, R. M., Friedewald, W. T., Schlesselman, J. J., DeMets, D. L., Ellenberg, J. H., Gail, M. H., & Ware, J. H. (1976). Randomized Clinical Trials: Perspectives on Some Recent Ideas. *The New England Journal of Medicine*, *295*(2), 74–80.

Callender, C., & Cohen, J. (2006). There Is No Special Problem About Scientific Representation. *Theoria: An International Journal for Theory, History and Foundations of Science*, *21*(1(55)), 67–85.

Carnap, R. (1953). On the Comparative Concept of Confirmation. *The British Journal for the Philosophy of Science*, *3*(12), 311–318.

Carnap, R. (1962). *Logical Foundations of Probability* (2nd ed.). The University of Chicago Press. (Original work published 1950)

Cartwright, N. (1979). Causal Laws and Effective Strategies. *Noûs*, *13*(4), 419–437.

Cartwright, N. (1999). *The Dappled World*. Cambridge University Press.

Cartwright, N. (2007). Are RCTs the Gold Standard? *BioSocieties*, *2*(1), 11–20.

Corcoran, J., & Hamid, I. S. (2016). Schema. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2016). Metaphysics Research Lab, Stanford University.

da Costa, N., & French, S. (2000). Models, Theories, and Structures: Thirty Years On. *Philosophy of Science*, *67*, S116–S127.

Davey, K. (2021). Inference to the Best Explanation and Norton's Material Theory of Induction. *Studies in History and Philosophy of Science Part A*, *85*, 137–144.

Day, T., & Kincaid, H. (1994). Putting Inference to the Best Explanation in its Place. *Synthese*, *98*(2), 271.

de Grefte, J. (2020). Epistemic Benefits of the Material Theory of Induction. *Studies in History and Philosophy of Science Part A*, *84*, 99–105.

Deaton, A., & Cartwright, N. (2018). Understanding and Misunderstanding Randomized Controlled Trials [Randomized Controlled Trials and Evidence-based Policy: A Multidisciplinary Dialogue]. *Social Science & Medicine*, *210*, 2–21.

Douglas, H. (2000). Inductive Risk and Values in Science. *Philosophy of Science*, *67*(4), 559–579.

Earman, J. (2020). Quantum Sidelights on the Material Theory of Induction. *Studies in History and Philosophy of Science Part A*, *82*, 9–16.

Earman, J., & Salmon, W. C. (1992). The Confirmation of Scientific Hypotheses. In M. H. Salmon, J. Earman, C. Glymour, J. G. Lennoz, P. Machamer, J. E. McGuire, J. D. Norton, W. C. Salmon, & K. F. Schaffner (Eds.), *Introduction to the Philosophy of Science*. Hackett Publishing Company.

Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science*, *66*, S362–S378.

Fitelson, B. (2001). *Studies in Bayesian Confirmation Theory* [Doctoral dissertation, University of Wisconsin-Madison].

Fitelson, B. (2006). Logical Foundations of Evidential Support. *Philosophy of Science*, *73*(5), 500–512.

Fitelson, B. (2007). Likelihoodism, Bayesianism, and Relational Confirmation. *Synthese*, *156*(3), 473–489.

Fitelson, B., & Sober, E. (2011). Favoring, Likelihoodism, and Bayesianism. *Philosophy and Phenomenological Research*, *83*(3), 666–672.

Frigg, R. (2010). Models and Fiction. *Synthese*, *172*(2), 251–268.

Frigg, R., & Nguyen, J. (2017). Models and Representation. In L. Magnani & T. Bertolotti (Eds.), *Springer Handbook of Model-Based Science* (pp. 49–102). Springer.

Fuller, J. (2019). The Confounding Question of Confounding Causes in Randomized Trials. *The British Journal for the Philosophy of Science*, *70*(3), 901–926.

Gelman, A., & Shalizi, C. R. (2013). Philosophy and the Practice of Bayesian Statistics. *British Journal of Mathematical and Statistical Psychology*, *66*(1), 8–38.

Gemes, K. (2007). Verisimilitude and Content. *Synthese*, *154*(2), 293–306.

Genta, B. S. (2020). How to Think About Analogical Inferences: A Reply to Norton. *Studies in History and Philosophy of Science Part A*, *82*, 17–24.

Giere, R. N. (2004). How Models Are Used to Represent Reality. *Philosophy of Science*, *71*(5), 742–752.

Glass, D. H., & McCartney, M. (2015). A New Argument for the Likelihood Ratio Measure of Confirmation. *Acta Analytica*, *30*(1), 59–65.

Godfrey-Smith, P. (2006). The Strategy of Model-Based Science. *Biology and Philosophy*, *21*(5), 725–740.

Goodman, N. (1983). *Fact, Fiction and Forecast* (4th ed.). Harvard University Press.

Graves, J. C. (1974). Uniformity and Induction. *The British Journal for the Philosophy of Science*, *25*(4), 301–318.

Hájek, A., & Joyce, J. M. (2008). Confirmation. In S. Psillos & M. Curd (Eds.), *Routledge Companion to the Philosophy of Science* (pp. 115–29). Routledge.

Harman, G. (1984). Logic and Reasoning. *Synthese*, *60*(1), 107–127.
Harman, G. (2002). Internal Critique: A Logic Is Not a Theory of Reasoning and a Theory of Reasoning Is Not a Logic. In D. M. Gabbay, R. H. Johnson, H. J. Ohlbach, & J. Woods (Eds.), *Studies in Logic and Practical Reasoning* (pp. 171–186, Vol. 1). Elsevier.
Harman, G. (2009). Field on the Normative Role of Logic. *Proceedings of the Aristotelian Society*, *109*, 333–335.
Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
Hawthorne, J. (2018). Inductive Logic. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2018). Metaphysics Research Lab, Stanford University.
Hempel, C. G. (1943). A Purely Syntactical Definition of Confirmation. *Journal of Symbolic Logic*, *8*(4), 122–143.
Hempel, C. G. (1945a). Studies in the Logic of Confirmation (I.) *Mind*, *54*(213), 1–26.
Hempel, C. G. (1945b). Studies in the Logic of Confirmation (II.) *Mind*, *54*(214), 97–121.
Henderson, L. (2014). Bayesianism and Inference to the Best Explanation. *British Journal for the Philosophy of Science*, *65*(4), 687–715.
Hilpinen, R. (1976). Approximate Truth and Truthlikeness. In M. Przelecki, K. Szaniawski, R. Wójcicki, & G. Malinowski (Eds.), *Formal Methods in the Methodology of the Empirical Sciences*. Reidel.
Hoefer, C. (2007). The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance. *Mind*, *116*(463), 549–596.
Howson, C. (2000). *Hume's Problem: Induction and the Justification of Belief*. Oxford University Press.
Hughes, R. I. G. (1997). Models and Representation. *Philosophy of Science*, *64*, S325–S336.
Iranzo, V. (2008). Bayesianism and Inference to the Best Explanation. *Theoria: An International Journal for Theory, History and Foundations of Science*, *23*(1(61)), 89–106.
Karanicolas, P. J., Farrokhyar, F., & Bhandari, M. (2010). Blinding: Who, What, When, Why, How? *Canadian Journal of Surgery. Journal Canadien de Chirurgie*, *53*(5), 345–348.
Kellert, S. H., Longino, H. E., & Waters, C. K. (2006). *Scientific Pluralism* (S. H. Kellert, H. E. Longino, & C. K. Waters, Eds.). University of Minnesota Press.
Khalifa, K., Millson, J., & Risjord, M. (2017). Inference to the Best Explanation: Fundamentalism's Failures. In K. McCain & T. Poston (Eds.), *Best Explanations: New Essays on Inference to the Best Explanation* (pp. 80–96). Oxford University Press.
Kolmogorov, A. N. (1950). *Foundations of the Theory of Probability*. Chelsea. (Original work published 1933)
Kuipers, T. A. (1987a). A Structuralist Approach to Truthlikeness. In T. A. Kuipers (Ed.), *What Is Closer-To-The-Truth? A Parade of Approaches to Truthlikeness* (pp. 79–99). Rodopi.
Kuipers, T. A. (1987b). *What Is Closer-To-The-Truth? A Parade of Approaches to Truthlikeness* (T. A. Kuipers, Ed.). Rodopi.
Kyburg, H. E. (1983). Recent Work in Inductive Logic. In T. Machan & K. Lucey (Eds.), *Recent Work in Philosophy* (pp. 87–150). Rowman & Allanheld.
Larroulet Philippi, C. (2022). There Is Cause to Randomize. *Philosophy of Science*, *89*(1), 152–170.
Levy, A. (2012). Models, Fictions, and Realism: Two Packages. *Philosophy of Science*, *79*(5), 738–748.

Levy, A. (2015). Modeling without Models. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *172*(3), 781–798.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability* (pp. 263–93). University of California Press.

Lipton, P. (2004). *Inference to the Best Explanation* (2nd ed.). Routledge.

Livengood, J., & Korman, D. Z. (2020). Debunking Material Induction. *Studies in History and Philosophy of Science Part A*, *84*, 20–27.

Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.

MacFarlane, J. (2004, April). *In What Sense (If Any) Is Logic Normative for Thought?* [Unpublished manuscript, Central Division APA 2004 presentation].

Magnani, L. (2009). *Abductive Cognition. The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*. Springer.

McCaskey, J. P. (2020). Reviving Material Theories of Induction. *Studies in History and Philosophy of Science. Part A*, *83*, 1–7.

McCaskey, J. P. (2022, February). *Making of the Problem: Induction from Socrates to Popper* [Unpublished manuscript].

Mill, J. S. (1974). *A System of Logic*. University of Toronto Press. (Original work published 1843)

Miller, D. (1974). Popper's Qualitative Theory of Verisimilitude. *The British Journal for the Philosophy of Science*, *25*(2), 166–177.

Miller, D. (1976). Verisimilitude Redeflated. *The British Journal for the Philosophy of Science*, *27*(4), 363–381.

Miller, D. (1978). The Distance Between Constituents. *Synthese*, *38*(2), 197.

Miller, D. (1994). *Critical Rationalism: A Restatement and Defence*. Open Court.

Miller, D. (2006). *Out Of Error: Further Essays on Critical Rationalism*. Routledge.

Milne, P. (1996). log[P(h/eb)/P(h/b)] Is the One True Measure of Confirmation. *Philosophy of Science*, *63*(1), 21–26.

Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2013). The Humble Bayesian: Model Checking From a Fully Bayesian Perspective. *British Journal of Mathematical and Statistical Psychology*, *66*(1), 68–75.

Morreau, M. (2010). It Simply Does Not Add Up: Trouble With Overall Similarity. *The Journal of Philosophy*, *107*(9), 469–490.

Newton, I. (1999). *The Principia: Mathematical Principles of Natural Philosophy* (B. Cohen & A. Whitman, Trans.). University of California Press. (Original work published 1723)

Neyman, J. (1955). The Problem of Inductive Inference. *Communications on Pure and Applied Mathematics*, *8*(1), 13–45.

Niiniluoto, I. (1998). Verisimilitude: The Third Period. *The British Journal for the Philosophy of Science*, *49*(1), 1–29.

Northcott, R. (2013). Verisimilitude: A Causal Approach. *Synthese*, *190*(9), 1471–1488.

Norton, J. D. (2003). A Material Theory of Induction. *Philosophy of Science*, *70*(4), 647–670.

Norton, J. D. (2005). A Little Survey of Induction. In P. Achinstein (Ed.), *Scientific Evidence: Philosophical Theories and Applications* (pp. 9–34). Johns Hopkins University Press.

Norton, J. D. (2008). The Dome: An Unexpectedly Simple Failure of Determinism. *Philosophy of Science*, *75*(5), 786–798.

Norton, J. D. (2010). There Are No Universal Rules for Induction. *Philosophy of Science*, *77*(5), 765–777.

Norton, J. D. (2011). History of Science and the Material Theory of Induction: Einstein's Quanta, Mercury's Perihelion. *European Journal for Philosophy of Science*, *1*(3), 3–27.

Norton, J. D. (2021a). Author's Responses. *Studies in History and Philosophy of Science Part A*, *85*, 114–126.

Norton, J. D. (2021b). *The Large-Scale Structure of Inductive Inference* (03/2021). Unpublished manuscript.

Norton, J. D. (2021c). *The Material Theory of Induction*. University of Calgary Press.

Oddie, G. (2016). Truthlikeness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016). Metaphysics Research Lab, Stanford University.

Ohnesorge, M. (2022). *Newton as Geodesist: The Problem of the Earth's Figure and the Argument for Universal Gravitation* [https://go.aps.org/geodesy].

Papineau, D. (1994). The Virtues of Randomization. *The British Journal for the Philosophy of Science*, *45*(2), 437–450.

Parker, W. S. (2020). Model Evaluation: An Adequacy-for-Purpose View. *Philosophy of Science*, *87*(3), 457–477.

Peden, W. (2019). Direct Inference in the Material Theory of Induction. *Philosophy of Science*, *86*(4), 672–695.

Peto, R., Pike, M. C., Armitage, P., Breslow, N. E., Cox, D. R., Howard, S. V., Mantel, N., McPherson, K., Peto, J., & Smith, P. G. (1976). Design and Analysis of Randomized Clinical Trials Requiring Prolonged Observation of Each Patient. I. Introduction and Design. *British Journal of Cancer*, *34*(6), 585–612.

Popper, K. R. (1963). *Conjectures and Refutations*. Routledge.

Quine, W. V. (1969). *Ontological Relativity and Other Essays*. Columbia University Press.

Reichenbach, H. (1949). *The Theory of Probability* (2nd). University of California Press.

Reiss, J. (2020). What Are the Drivers of Induction? Towards a Material Theory⁺. *Studies in History and Philosophy of Science. Part A*, *83*, 8–16.

Roche, W., & Shogenji, T. (2014). Dwindling Confirmation. *Philosophy of Science*, *81*(1), 114–137.

Roche, W., & Sober, E. (2013). Explanatoriness Is Evidentially Irrelevant, or Inference to the Best Explanation Meets Bayesian Confirmation Theory. *Analysis*, *73*(4), 659–668.

Savage, L. J. (1972). *The Foundations of Statistics* (2nd ed.). Dover Publications, Inc.

Schurz, G. (2019). *Hume's Problem Solved: The Optimiality of Meta-Induction*. The MIT Press.

Schwartz, D., Flamant, R., & Lellouch, J. (1980). *Clinical Trials*. Academic Press.

Sellars, W. (1953). Inference and Meaning. *Mind*, *62*(247), 313–338.

Sober, E. (2008). *Evidence and Evolution: The Logic Behind the Science*. Cambridge University Press.

Sprenger, J., & Hartmann, S. (2019). *Bayesian Philosophy of Science*. Oxford University Press.

Stegenga, J. (2015). Herding QATs: Quality Assessment Tools for Evidence in Medicine. In P. Huneman, G. Lambert, & M. Silberstein (Eds.), *Classification, Disease and Evidence: New Essays in the Philosophy of Medicine* (pp. 193–211). Springer.

Steinberger, F. (2022). The Normative Status of Logic. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2022). Metaphysics Research Lab, Stanford University.

Suárez, M. (2004). An Inferential Conception of Scientific Representation. *Philosophy of Science*, *71*(5), 767–779.

Suppes, P. (1960). A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences. *Synthese*, *12*(2/3), 287–301.

Suppes, P. (1978). The Plurality of Science. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1978*, 3–16.

Tichý, P. (1974). On Popper's Definitions of Verisimilitude. *The British Journal for the Philosophy of Science*, *25*(2), 155–160.

Ubbink, J. B. (1960). Model, Description and Knowledge. *Synthese*, *12*(2/3), 302–319.

Urbach, P. (1985). Randomization and the Design of Experiments. *Philosophy of Science*, *52*(2), 256–273.

Vassend, O. B. (2020). A Verisimilitude Framework for Inductive Inference, with an Application to Phylogenetics. *The British Journal for the Philosophy of Science*, *71*(4), 1359–1383.

Worrall, J. (2002). What Evidence in Evidence-Based Medicine? *Philosophy of Science*, *69*(S3), S316–S330.

Worrall, J. (2007). Evidence in Medicine and Evidence-Based Medicine. *Philosophy Compass*, *2*(6), 981–1022.

Yusuf, S., Held, P., Teo, K. K., & Toretsky, E. R. (1990). Selection of Patients for Randomized Controlled Trials: Implications of Wide or Narrow Eligibility Criteria. *Statistics in Medicine*, *9*(1-2), 73–83.

Zalabardo, J. (2009). An Argument for the Likelihood-Ratio Measure of Confirmation. *Analysis*, *69*(4), 630–635.