Probability and Inductive Logic

Antony Eagle

University of Adelaide

antonyeagle.org

2023-03-24

Abstract

Reasoning from inconclusive evidence, or ‘induction’, is central to science and any applications we make of it. For that reason alone it demands the attention of philosophers of science. This element explores the prospects of using probability theory to provide an inductive logic, a framework for representing evidential support. Constraints on the ideal evaluation of hypotheses suggest that overall support for a hypothesis is represented by its probability in light of the total evidence, and incremental support, or confirmation, indicated by an increased probability given the evidence than otherwise. This proposal is shown to have the capacity to reconstruct many canons of the scientific method and inductive inference. Along the way, significant objections are discussed, such as the challenge from inductive scepticism and the objection that the probabilistic approach makes evidential support arbitrary.

Contents

About this Book

Open-mindedness is a cognitive virtue – perhaps even a moral one (Arpaly 2011). To be open-minded is to be disposed to consider other views, avoiding rigidity in opinion. Yet it is not to be intellectually fickle. The open-minded person changes their mind in response to evidence, when confronted with new facts or a new interpretation of old facts. Open-mindedness is a virtue, while fickleness is a vice, because the former conduces to our cognitive aims – truth, knowledge, wisdom – while the latter does not. To put virtue into practice however requires understanding what it is to *respond to evidence*. Deductive logic can provide some background structure, but most of the time managing our states of opinion involves 'weighing evidence and judging probability' (Lipton 2004: 5). This book is about those notions.

In Section 1, I introduce our topic, distinguishing inductive inference from deductive logic. I proposing that 'inductive logic' could serve as a good label for the study of relations of evidential support and confirmation at a time; however, the connection to inference is more fraught. Section 2 embarks on the project of inductive logic by examining some previous qualitative approaches, attempting to use logical features of the evidence and hypothesis under consideration to ground relations of evidential support. I note some technical challenges for these proposals, but also note their susceptibility to Hume's venerable 'problem of induction'. Section 3 starts on the positive project, proposing to understand the degree to which a hypothesis is supported as relative to a particular ideal perspective on the hypothesis, given a background body of total evidence and a conception of evidential support. It is there argued that these epistemic perspectives must have a probabilistic structure, and that they are not necessarily to be identified with the actual attitudes of any individual. Thus the account on offer is broadly Bayesian, yet not wholly subjectivist. Having introduced the notion of overall degree of support, I turn in Section 4 to the notion of incremental confirmation of a hypothesis by evidence, sketching ways in which the Bayesian approach explains successes of the scientific method, improves on the qualitative accounts discussed in Section 2 and on other probability-based approaches, and ultimately provides a framework for induction. I respond there to some challenges to my proposal, while devoting the whole of Section 5 to what I see as the most significant obstacle facing any broadly Bayesian approach to evidential support, the problem of identifying and justifying an appropriate epistemic perspective to take on a given issue of theory choice.

This book sits in a long tradition of other works on broadly Bayesian approaches to epistemology and general philosophy of science.[1] Because of space constraints, interesting topics and challenges do not receive the attention they perhaps deserve in this book; the interested reader is encouraged to explore other treatments of the issues. In particular, I try to focus on the positive development of a defensible inductive logic; many thoughtful objections receive no overt response here, though I believe my approach is set up in a way that forestalls some quite prominent criticisms. My overall approach displays the clear influence of Howson (2000) and Fitelson (2005), and I draw in part on my own earlier work (Eagle 2011; Eagle 2016a), though I have changed my mind on many things. I have silently made notation in quotations uniform with the remainder of the book. I am grateful for the forbearance of the series editors as I worked to complete this book, to students in my honours seminar on Probability and Inductive Logic in 2021 who endured an early prototype, to Marshall Abrams for facilitating a 2022 PSA symposium on randomness where a version of §5.5 was presented, to Coffee in Common and Crave Specialty Coffee for providing congenial working environments, and to Lizzie, Sylvester, and Jonquil for their sometimes stretched patience and constant encouragement.

[1] Even a non-exhaustive list is exhausting (Carnap 1962; Horwich 1982; Earman 1992; Howson and Urbach 1993; Skyrms 2000; Hacking 2001; Bovens and Hartmann 2003; Strevens 2006; Jeffrey 2008; Norton 2011; Weisberg 2011; Bradley 2015; Schupbach 2022; Titelbaum 2022a; b). That is without including works on the philosophy of probability more broadly.

Section 1

Logic, Induction, and Inductive Logic

In this opening section, I introduce the notion of inductive logic. Given the plurality of ways that philosophers have approached logic and induction, it is important to get clear on these topics first in order to understand how they may delineate the field of inductive logic. I begin in §1.1 with a discussion of logic as a formal relation among sentences, contrasting that in §1.2 with induction characterised as a species of ampliative reasoning. Meshing consequence relations at a time with norms of reasoning over time is at first glance an incoherent project, and this poses an initial challenge to the project of inductive logic (§1.3). I suggest in §1.4 that this challenge can be defused by focusing on an inductive logic that is not a theory of inductive reasoning, but is rather a synchronic theory of evidential support and confirmation.

1.1 Logic

If Jonquil is younger than Sylvester, Sylvester is older than Jonquil. Jonquil is younger than Sylvester. It is a **formal consequence** of these first two sentences that Sylvester is older than Jonquil. Logic is the science of formal consequence.

Formal consequence is a relation among sentences that obtains in virtue of the form of those sentences. The form of a sentence is not merely a product of its constituent letters, but also its syntactic structure. Accordingly, the form of a given sentence (string of symbols) will depend also on a choice of language, which will determine the phrasal structure assigned to the sentence. (Or perhaps, each formal language has a particular approach to modelling the syntactic structure of a natural language sentence.) The most common languages introduced in elementary logic courses offer quite different analyses of the structure of our opening sentence *If Jonquil is younger than Sylvester, then Sylvester is older than Jonquil*. Sentential logic offers an analysis into 'basic' sentences, which are connected by the structural word *if*; it says the sentence has the form $(J \rightarrow S)$. Predicate logic goes further; one might fruitfully analyse those basic sentences, perhaps yielding the form $(Yjs \rightarrow Osj)$ involving names and two binary relational predicates. Other languages may go even further, might be itself analysed — perhaps the comparative form *is older than* is to be analysed morphologically, and that could be reflected in a conception of structure adopted by some logical language.

A language not only delineates the form of sentences, but also assigns meanings to the expressions in the language. A sentence ϕ is a formal consequence of some sentences Γ, relative to some language, just in case there is no possible way to uniformly reinterpret the non-structural expressions in Γ and ϕ so as to make all the sentences in Γ true and ϕ false (Tarski 1983: 416–17). So S is a formal consequence of *If* J, S and J in classical sentential logic because, keeping fixed the meaning that language assigns to *if*, whatever meanings are assigned to J and S will ensure the truth of S given the truth of *If* J, S and J.

Formal consequence should be distinguished from other notions of consequence – for example, causal consequence ('the vase broke as a consequence of my knocking it over'). Causal consequence is contingent and substantive. Whether a causal relation holds between two events can vary with the background conditions (such as the presence of pre-empting causes) and the laws of nature. Formal consequence, while relative to a language, is an internal relation, necessitated by the forms of the sentences involved and the meanings assigned by that language.

Logic on this conception is *not* a theory of reasoning or inference (Harman 1986: 3–6). This can be obscured by some otherwise attractive proof systems for logic, such as natural deduction, which tend to encourage a conflation between the principles describing the acceptable structures for formal proofs, and principles of natural inference. Yet while formal consequence is a relation between sentences in a language at one time (a **synchronic** relation), reasoning is a relation between states of opinion over time (a **diachronic** relation). Given that inference essentially relates states of opinion, a correct theory of reasoning will depend on what is distinctive about such states that may not be shared with more general quasi-linguistic entities like sets of sentences or propositions. (Though sometimes a state of opinion might be represented by a set of sentences, but even there the question which sets of sentences can represent a state of opinion is not trivial.)

Perhaps facts about consequence in a speaker's language can be pertinent to evaluating inferences that a speaker might make while reasoning. For example, perhaps the fact that a speaker has a belief in a sentence with the structure $(\phi \rightarrow \psi)$ and believes ϕ might give that speaker *prima facie* grounds to infer ψ. But even this holds only defeasibly. What is plausible is that there is something synchronically unreasonable about having a state of opinion in which someone is committed to Γ and $\neg\phi$, where ϕ is a formal consequence of Γ. But even that might not be all-things-considered unreasonable; what if the believer doesn't recognise the inconsistency (Harman 1986: 17–19)?

1.2 Induction

Induction, by contrast, is generally characterized as a species of inference. Howson characterises the central question of induction as being how 'our hard-won factual knowledge' iss able to be 'secured by any process of demonstrably sound reasoning' (2000: 1). Inductive reasoning is the paradigm of scientific reasoning, as when the scientific community adopted Hutton and Lyell's 'Uniformitarian' approach to geology (Lyell 1830). The Uniformitarians argued that the same geological processes (erosion, deposition, lithification, orogeny, etc.,) active today are also those responsible for shaping the landscape of the Earth throughout time. Ripple marks in an exposed vertical rock layer, far from the sea, are an apparent mystery. The uniformitarian explains them as the concatenation over geological time of familiar processes: postulating an ancient shallow sea in which the ripple marks were formed, the covering of those rippled layers by subsequent sediments, the gradual folding and uplift of the resulting sedimentary rock into a mountain range, and the erosion of that range over millions of years to reveal the ripples again to the eye (Drexel, Preiss, and Parker 1993: 171–97).[1] This simple and elegant hypothesis was opposed to the many unattested geological mechanisms invoked *ad hoc* by the rival 'Catastrophist' school. The inductive inference here was to one of these rival hypotheses as the **best explanation** of the agreed geological data (Lipton 2004). The preferred Uniformitarian explanation postulates a continuity or resemblance between present and past geological processes. In postulating this continuity, the Uniformitarians also endorse another canonical form of inductive inference, from evidence of past geology to predictions concerning future geology (Hume 1777). This itself is an example of a slightly more general but very common pattern of reasoning, **inverse inference** from features of a sample to those of the population from which it is drawn; which in turn perhaps ought to be simply unified with inference to the best explanation (Harman 1965). I will treat all of these as kinds of inductive inference.

Inference to the best explanation and inverse inference are both exemplified in ordinary reasoning as well as scientific contexts. We conclude that the last evening train from the will depart late, on the basis that it's always been late; we may also conclude in turn that something is a little different about the passengers on the last train, perhaps hypothesising that they are particularly prone to the sorts of behaviours that cause delays.

While there are interesting special features of inverse inference as opposed to the inference to the best explanation more generally, both exhibit an 'ampliative' character that reasoning constrained to mimic formal consequence does not. In the example above, the Uniformitarian hypothesis is not a formal consequence of the given geological evidence, because Catastrophism is consistent with the evidence but incompatible with Uniformitarianism. The resources provided by formal consequence allow the exclusion of hypotheses that are inconsistent with the evidence, but do not provide grounds to favour any of the hypotheses consistent with the evidence over any other. By contrast, reasonable inductive inference can lead us to favour one of many coherent hypothesis, as in the examples above. It is coherent to suppose the future quite unlike the past; yet we habitually infer that the fu-

[1] This can include sudden changes, such as the radical reshaping of the 'Channeled Scablands' of western Washington by flooding (Baker 2009), as long as those changes themselves are examples of currently active geological processes.

ture will broadly resemble the past. As such, inductive inference 'goes beyond' the evidence: the result of inductive inference is not entailed by the logical content of the evidence that prompted the inference. It is that feature which will be our principal concern in what follows.

Here, and throughout, evidence is treated as **propositional**. It is the evidence *that volcanic rocks intruded into older sedimentary rocks* that was grounds for Hutton to favour the Uniformitarian idea that the formation and erosion of rock was a continuous and ongoing process, over the 'Neptunist' idea that all rock was precipitated from a primordial sea in an event which has no current parallel. The visible rocks are thus not the evidence, though the evidence is about them, and is acquired as a response to engaging with them. Likewise, when it is said that smoke is evidence of fire, this must be understood as elliptical for something like *that smoke is present is evidence for the presence of fire*.

Hume's famous discussion of induction is largely prompted by his recognition that inductive inference isn't exhausted by formal consequence. In section IV of the *Enquiry* Hume is concerned with discovering the foundation for our 'inferences from experience' (Hume 1777: ¶4.21), and suggests that

> all arguments from experience are founded on the similarity, which we discover among natural objects, and by which we are induced[2] to expect effects similar to those, which we have found to follow from such objects. (Hume 1777: ¶4.19)

Here induction is essentially diachronic. Of course Hume has a particular conception of the largely unconscious psychological processes by which we are led to beliefs that 'go beyond the evidence of our memory and senses' (Hume 1777: ¶4.4). This naturalistic account of induction leads Hume to conclude that the interesting ideas about inductive reasoning principally concern how it happens, as opposed to attempts to provide a rational justification for that reasoning, which he thinks are not available. (We return to Hume's inductive scepticism below, §2.4.)

Moreover, induction involves *substantive* reasoning – it is not merely formal in character. A successful inductive inference is grounded not in

[2] Strikingly, this is the only word cognate with 'induction' appearing in Hume's text, where it appears to denote an involuntary process forming our expectations.

the formal structure of the evidence, but in its content. As Norton says, 'inductions ultimately derive their licenses from facts pertinent to the matter of the induction' (Norton 2003: 650). Hume recognised that our inductive habits rested on the pattern of causal relationships evident in our experience, and the 'supposition' that future patterns 'will be conformable to the past' (Hume 1777: ¶4.19). What it will be reasonable to conclude will therefore depend crucially on what kinds of patterns of causal relationships are exemplified in your evidence; and one's inductive reasoning leads to inductive knowledge only if that evidence is itself conformable to the true system of causal relations instantiated in one's reality. We know that the geological future of the Himalayas will involve the continued uplift of the ranges as the Indian plate ploughs into the Tibetan plateau. Were the geological laws and mechanisms different, this conclusion (even if true) might not amount to knowledge; perhaps the uplift of the Himalayas could have had some catastrophist explanation instead.

1.3 Inductive Logic

Induction is about diachronic substantive reasoning from one state of opinion to conclusions that go beyond our experience, guided by the impact of new evidence. Logic is about synchronic formal coherence relations among sentences and sets of sentences. It might thus seem obvious that any project of **inductive logic** rests on a mistake:

> if we clearly distinguish reasoned change in view from argument, we cannot suppose that the existence of inductive reasoning by itself shows there is such a thing as inductive argument, nor can we suppose that it shows there is an inductive logic. (Harman 1986: 5)

We ought not hastily follow Harman's hint and embrace scepticism about inductive logic. Deductive inference takes us from one state of opinion to another. While that transition is not governed by logic, it may well be that logical properties of those initial and final states are part of the account of which transitions are rationally permissible. 'Inductive logic' might be an apt name for the generalization of this project that involves characterizing synchronic constraints that are needed to understand diachronic inductive inference.

To elaborate, consider the role logic plays in justifying inference from premise to consequence. Let us represent a state of opinion as a set of sentences, for illustration. Suppose I begin with an opinion including $(\phi \to \psi)$, and acquire the new evidence ϕ. This new evidence is added, unreasoningly and automatically, to my existing state of opinion. But its arrival does prompt me to reason, to revise the newly augmented opinion, and I do so by reasoning, coming to adopt a belief that ψ. This inference isn't obligatory in light of my existing opinion – I could give up my existing conditional opinion instead. But some logical facts do inform it. For example, my existing opinion was *lacking* ψ, by which I mean that its constituents, including the new evidence, entailed ψ but it did not contain ψ. That a state of opinion is lacking with respect to some sentence is a synchronic logical property of that state of opinion. The state to which I reasoned was not lacking in this sense, and it is natural to think that when I reasoned to remedy a lack in my state of opinion, I was at least doing something rationally permissible. Thus we might use a logical constraint to articulate a proposed diachronic norm: if a given state of opinion is lacking with respect to some claim, it is always rationally permissible to reason to a subsequent state of opinion which is not lacking in this respect. This norm will vindicate my reasoning as rationally permissible, though of course the state of opinion which would have resulted had I simply abandoned my view that ϕ would also have not been lacking with respect to ψ, and hence my reasoning was not rationally obligatory.

In this proposal, a formal logical property provides an important constraint on answers to a normative question about synchronic belief, namely, *what can we rationally believe given a certain body of evidence?*. It would not be the whole story of rational inference, because it offers no theory of how one ought to revise a coherent state on receipt of new evidence. So the synchronic question is distinct from, and prior to, the diachronic question *what ought we come to believe given some newly acquired evidence?* That second question genuinely concerns inference, the rationality of changes of opinion. The first question concerns the coherence or structure, at a time, of a state of opinion which already incorporates some body of evidence. Insofar as they are both normative, the two questions should not be wholly independent of one another, for one may hope that any permissible change of opinion will result in a permissible state of opinion.

Nevertheless, a purely synchronic constraint on our opinions can be combined with many different accounts of how they can be combined over time. One might have a very liberal view, that all that matters is coherence at each moment; it may not be irrational to have radical discontinuities between successive coherent states of opinion (Hedden 2015: 6–9, 28–55). An update rule given such liberality might simply say: if one receives the new information that p, update to a coherent state that includes p. One might require something slightly stronger: if one's attitude doesn't include one's total evidence, one is irrational; this will requires that new information is always included, alongside whatever elements of the old are retained. One might have various stricter views that require stronger connections between existing and revised states of opinion. For example, the AGM theory of belief revision requires that a revised state must be a 'minimal revision' sufficient to incorporate the new information p (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988). On this theory, to update one first removes any commitment to $\neg p$ from a state of opinion, but without removing any beliefs unnecessarily, then expanding that amended state by p. Another quite different theory of belief revision requires not that successive states of opinion are as close to one another as possible, given coherence, but that successive states of opinion must agree with one another on 'structural' principles. For example, if doxastic states are represented by probability assignments (a view we will discuss throughout this book), the updating rule Conditionalization allows quite radical changes in unconditional probability assignments, but requires 'rigidity': agreement on conditional probabilities between existing and revised assignments (Weisberg 2009: 14–16).

Let us set aside the issue of diachronic constraints on rationality. The suggestion we will explore is that 'inductive logic' would be an apt label for the attempt to generalize the role of synchronic features of states of opinion in understanding reasoning beyond purely logical features. This is a fraught topic. There is some controversy over the principles characterizing formal features of a set of sentences – for example, are they the principles of classical logic, or perhaps some non-classical relevance logic? But these controversies are dwarfed in prevalence and scope by disputes over whether a given inductive opinion is licensed by some evidence. As cases of scientific controversy illustrate, even the best positioned epistemic agents can disagree over what states of opinion are appropriate, even when they incorporate a shared body of evidence. Supposing we agree that logical consistency is necessary for a state of opinion to be rationally permitted, is it possible to articulate any plausible characterization of those doxastic states that are rationally permissible given some body of evidence? This is the central question of inductive logic.

1.4 Evidential Support

Suppose we have a piece of inductive reasoning; someone comes to adopt the belief that Labor will win the next general election, prompted to that reasoning by some newly acquired evidence about polling. This is an example of inverse inference. The prior state includes some general background beliefs about how the election works, information that makes the statistical properties of the whole voting population determinative of the election result. To that background opinion is added some new evidence about Labor's dominance in the polls, plus doubtless some information indicative of the representativeness of the polls. This acquisition is an augmentation of the background opinion, though again, it is not itself taken to be a reasoned change of opinion, but is assumed to be automatic. The information in this augmented existing state of opinion does not entail that Labor will win the election, so the state of opinion is not lacking with respect to that claim (unlike the example from §1.3). But we might think there is a kind of attenuated lack here: while the existing opinions don't entail that Labor will win, they seem to strongly suggest that Labor will win. This 'suggestiveness' of the evidence, at least relative to the existing state of opinion, is a synchronic feature of that augmented state. That the existing state of opinion fails to include a claim that is suggested, on balance, by the other claims it does include is another synchronic feature, analogous to a state being found wanting with respect to a given claim. Let's say that a state of opinion which fails to include some ϕ it suggests is *unimpressed* with respect to ϕ. Then we can propose various diachronic norms on reasoning; for example, we could say that it is always rationally permissible to infer to a state of opinion which is not unimpressed with respect to ϕ (i.e., a state that includes ϕ, or which fails to include some of the background opinion which suggested ϕ) from a state which was unimpressed with respect to ϕ, and so on.

I don't want to go into further detail concerning the above suggestion, which is merely illustrative (and ultimately, again, to continue engaging the issue of diachronic norms would take us too far afield). What I am interested in is this notion of *suggestiveness*. To coin an official label, let's say that this is a case where the polling evidence provides *confirmation* of the hypothesis that Labor will win. In fact, that piece of merely potential evidence can be recognised as supportive of a Labor victory, given the background opinion, before it is even obtained. It is because it is a supportive piece of evidence that, when it is eventually gathered, it forces a change of opinion. But whether a relation of evidential support obtains is supposed to be a synchronic feature of a given state of opinion, one that is prior to, but obviously constrains, appropriate reasoning from that state. The obvious proposal is that the theory of such a relation would be aptly termed an *inductive logic*.

That a doxastic attitude is rational only if it is based on all the evidence available to an agent has been termed the Principle of Total Evidence (Carnap 1947: 138–39). The relation of confirmation just introduced relates a single piece of evidence to a hypotheses, and hence can't justify a doxastic attitude to a hypothesis but itself. There is a relation of *evidential support* that holds between a body of (actual or potential) evidence and a hypothesis, not a single piece of evidence in isolation. In the election example above, the background information was initially distinguished from the newly acquired polling data. But to understand the way that the polling data is supportive of a Labor victory, in a way that figures in belief, we must include both in the body of evidence under consideration. To the extent that it is coherent to wonder how a single piece of evidence bears on a hypothesis, it will be in virtue of counterfactually supposing that single piece of evidence to be the total evidence and noting what it supports – typically, it will provide very little evidential guidance in isolation from any auxiliary assumptions (Quine 1951).

We came to the notion of evidential support by noting that in good inductive inferences, the premises support the conclusion, given the background assumptions in play. Another reason for distinguishing the diachronic and synchronic aspects of inductive inference is that there are cases where an inductive inference isn't licensed, but relations of evidential support are still apparently present. For example, if we have small sample of people relative to a whole population, any evidence that all smokers in the sample have impaired lung capacity will not be sufficient reason to accept that all smokers have impaired lung capacity. Nevertheless, the evidence of that sample does seem to confirm that conclusion, even if it is not enough to reach the level where an inference is warranted.[3]

On this conception, part of the study of inductive inference would concern the right way to characterise confirmation and evidential support; this would be inductive logic proper.[4] The remainder would concern how facts about evidential support feed into principles of rational belief revision, such as the proposal just above, or, as another example, the principle that when some evidence E supports H, that is a reason for you to believe H.[5]

There is no guarantee that factoring the study of induction into a synchronic and a diachronic component will bear fruit. Perhaps the difficulty of articulating principles of diachronic rationality – telling us how to take account of relations of evidential support in our reasoning – indicates that we should approach inductive reasoning in some different way; but perhaps not. The only way to proceed here is to attempt to construct a theory of evidential support, and evaluate the project of inductive logic in light of the fruits or failures of that attempt. That, in any case, is the constructive spirit behind my approach in this book.

3 Forcing this example into the Procrustean bed of significance testing (Fisher 1935), we might say that the sample statistics are suggestive without being 'significant', where the latter is a technical term for a result that (for devotees of the framework) prompts a change in view, namely, the rejection of the null hypothesis. It is a weakness of the significance testing paradigm that it has much to say about the notion of significance, and decisions to accept or reject hypothesis, and nothing at all to say about the prior notion of evidential support.

4 This is what Strevens (2004) calls an 'inductive framework'; he wants to reserve 'inductive logic' for a full account of abductive reasoning, a mistake in light of our discussion above.

5 This would again be a defeasible principle; one can imagine a circumstance in which you acquire evidence which supports H, but only to a small degree. For example, suppose you have some clinical trial data which shows a significant but small effect of a new painkiller. (Perhaps in the treatment group there is a small decrease in self-reported pain.) This is some sense a reason to believe the drug works. But now suppose you were the drug's discoverer, and a big believer in the efficacy of the drug. The surprisingly marginal results might prompt you to be *less* committed to the hypothesis that the drug works than you were. So while the evidence supports the hypothesis, it does not do so to an extent that licences your prior level of confidence, and hence you might be rational to take the supporting evidence as a reason to move to a subsequent state of opinion that was less committal about the hypothesis.

Section 2

Logic and Evidential Support

In this section I look at earlier approaches to the logic of induction, starting with the purely logical approach of Hypothetico-Deductivism in §2.1. Failures of that approach motivated Hempel's attempt to articulate necessary conditions on evidential support (§2.2), though they fell far short of sufficiency. Problems with Hempel's conditions reveal in §2.3 a more thoroughgoing difficulty with any broadly formal approach to evidential support, namely, the essential role of background assumptions. In §2.4 I suggest that background assumptions are also vital in a successful treatment of Hume's challenges to induction, and provide a viable way for inductive logic to thread the needle between scepticism and implausibility.

2.1 The Hypothetico-Deductive Approach

An admirably simple proposal would be to give a theory of evidential support using logical resources alone. That would give us an inductive logic both in my more analogical sense, and in a stricter sense, being more or less continuous with logic proper. One such proposal is the hypothetico-deductive theory of evidential support, which amounts to the thought that if some possible event is predicted by a hypothesis, then the evidence of that event's occurrence supports the hypothesis.

Consider the example from §1.4 involving a sample in which all smokers have impaired lung capacity. The idea was that this sample supports a hypothesis about the whole population. If the whole population is such that all smokers have impaired lung function, then one would predict (because it follows logically) that every sample of the population would exhibit this uniformity. Having in hand a sample that exhibits the uniformity and is just what we'd expect if the population were uniform, we get some partial support for the hypothesis that the whole population is uniform. This is made official in what has come to be called:

Nicod's Condition A universal generalisation *that all Fs are Gs* is supported by any instance of *an F which is a G*, and is undermined by any instance of *an F which is not a G*.(Hempel 1945a: 10)

Nicod's condition isn't a principle of logic: it links evidential support provided by E for H to an entailment relation from H to E, when H takes the form of a universal generalisation (given the auxiliary premise that there are some Fs). That this is a case of evidential support is supported by our judgements about cases.

Another principle of the same sort, again indicating the presence of evidential support given some logical facts, systematises the idea that evidential support should have deductive consequence as a limit case:

Entailment Condition If some evidence E entails a distinct non-tautological statement H, then E supports H. (Hempel 1945b: 103)

Alongside these sorts of principles, which indicate conditions which are *sufficient* for evidential support, there are also principles which indicate *necessary* conditions on evidential support. These principles state the logical consequences of an evidential support relation between some sentences. For example (noting that Hempel uses 'confirms' to refer to what we've been calling evidential support):

Equivalence Condition 'If an observation report confirms a hypothesis H, then it also confirms every hypothesis which is logically equivalent with H.' (Hempel 1945b: 103)

Putting the Entailment Condition together with Nicod's Condition leads to a problem however, Hempel's 'paradox of the ravens':

> if a is both a raven and black, then a certainly confirms S_1: '$\forall x(\text{Raven}(x) \rightarrow \text{Black}(x))$'; and if d is neither black nor a raven, d certainly confirms S_2: '$\forall x(\neg\text{Black}(x)) \rightarrow \neg\text{Raven}(x)$'. Let us now combine this simple stipulation with the Equivalence Condition: Since S_1 and S_2 are equivalent, d is confirming also for S_1; and thus, we have to recognize as confirming for S_1 any object which is neither black nor a raven. Consequently, any red pencil, any green leaf, and yellow cow, etc., becomes confirming evidence for the hypothesis that all ravens are black. (Hempel 1945a: 14, notation modified)

This is not plausible. Our judgment grows 'out of the feeling that the hypothesis that all ravens are black is about ravens, and not about non-black things, nor about all things'(Hempel 1945a: 17). The conclusion Hempel draws is that Nicod's Condition should be rejected in light of its 'several deficiencies' (Hempel 1945a: 22). Others have gone on to point out apparent direct counter-examples to the condition:

> The hypothesis under examination is 'All grasshoppers are located outside the County of Yorkshire'. The observation of a grasshopper just beyond the county border is an instance of this generalisation and, according to Nicod, confirms it. But it might be more reasonably argued that since there are no border controls or other obstacles restricting the movement of grasshoppers in that area, the observation of one on the edge of the county increases the probability that others have actually entered and hence undermines the hypothesis. ... this is a case where, relative to background information, the probability of some datum is reduced by a hypothesis ... which is therefore disconfirmed....(Howson and Urbach 1993: 129; see also Good 1961; Swinburne 1971: 326)

2.2 Hempel's Conditions of Adequacy

Hempel proposed a different approach. He listed 'conditions of adequacy' that any account of evidential support, or confirmation, had to meet; we could think of these as platitudes about evidential support, and specifically, the relationship between evidential support and logical consequence. One of these was the Entailment Condition. Another was the Consistency Condition, that some evidence is compatible with the set containing every hypothesis it supports (and hence that all the hypotheses supported by some evidence are consistent with one another). Hempel also endorsed the Equivalence Condition, actually taken to follow from this stronger principle:

Consequence Condition 'If an observation report confirms every one of a class K of sentences, then it also confirms any sentence which is a logical consequence of K.' (Hempel 1945b: 103)

Hempel rejected any vestige of the hypothetico-deductive approach, however. The most minimal version of the idea that evidential support is the 'converse' of prediction is given by this principle:

Converse Consequence Condition If some evidence E supports a hypothesis H, then E also supports any logically stronger hypothesis H' that entails H (Hempel 1945b: 104)

The Converse Consequence Condition might appear to have the support of intuition. Often, when we accept that some data supports a hypothesis, and later a stronger theory comes along that subsumes the first hypothesis, then the support provided by the data should carry over to the subsuming hypothesis too. For example, Darwin's theory of evolution was supported by the observational evidence of the diversity of species, so when Darwinian evolution was subsumed into the 'modern synthesis' together with Mendel's theory of inheritance, that observational evidence supported the modern synthesis too. That historical example looks like it fits the pattern of the Converse Consequence Condition.[1]

But in the presence of Hempel's other conditions of adequacy – indeed, even if the presence of just the Entailment Condition – the Converse Consequence Condition trivialises evidential support. Suppose E and H are any

[1] Nicod's Condition follows from Converse Consequence: $Fa \wedge Ga$ entails $Fa \rightarrow Ga$, so by the Entailment Condition, $Fa \wedge Ga$ supports $Fa \rightarrow Ga$; $\forall x(Fx \rightarrow Gx)$ entails $Fa \rightarrow Ga$; so by Converse Consequence, $Fa \wedge Ga$ supports $\forall x(Fx \rightarrow Gx)$.

two non-tautological claims, and that their disjunction $E \lor H$ is also non-tautological. Then we can show that E supports H (Morvan 1999; Moretti 2003: 299):

(1) E entails $E \lor H$ (logic)
(2) E supports $E \lor H$ (Entailment Condition, 1)
(3) H entails $E \lor H$ (logic)
(4) E supports H (Converse Consequence, 2, 3)

Hempel thus explicitly disavows the Converse Consequence Condition along with Nicod's Condition. The remaining conditions of adequacy are however far too weak to pin down any definite notion of evidential support. Hempel offered a theory of 'instance confirmation' intended to supplement the conditions of adequacy, but it has not been easy to make precise what he had in mind (Earman 1992: 66ff).

Indeed, we may not even want a notion of evidential support that coheres with these conditions. Hempel's instance confirmation was intended to capture the same intuition that prompted Nicod's Condition: that, at least sometimes, an instance of a generalization supports that generalization. If this intuition is to be respected, evidential support must be a relative, rather than an absolute notion. This is reflected in my deliberate choice of terminology, because a hypothesis can receive some support despite being on balance not credible. (To say that a hypothesis is 'confirmed', on the other hand, does seem to suggest that it is on balance credible.) But then we can envisage cases where our judgment is pretty straightforward that the conditions of adequacy are violated. Consider this case

Missing Bushwalker A bushwalker is missing. We know from his trip intentions that he's in the national park, but beyond that we have no reason to favour any hypothesis: he could be north or south of the river running through the middle of the park, and (independently) in either the east, central, or west blocks, giving 6 possible sectors to search (i.e., north-east, south-central, etc.). Consider the hypothesis W that he's in the west sector, and the hypothesis not-E that he's not in the east sector. Here W entails not-E. We now get a report from searchers on the ground: he's not in north-central.

It is plausible to think that this report supports W; both ways for W to be true remain live, and a non-W possibility has been eliminated. But it's equally plausible to think the report supports E and undermines not-E: both ways for E to be true remain live, and one of the non-E possibilities has been eliminated. This example violates both the Consistency and Consequence Conditions. The report supports a hypothesis W and doesn't support a consequence of it, not-E. The report supports E and W, even though they are inconsistent with one another.

The traditional wisdom concerning Hempel's conditions of adequacy is that Hempel himself is being pulled in two different directions by his prior concept of support (Huber 2008; but see Carnap 1962 for a different proposal regarding Hempel's confusion). On the one hand, he is pulled towards an absolute notion, on which evidence supports hypotheses that it makes absolutely plausible. On the other hand, his ideas about instance confirmation pull him towards a relative or incremental notion, on which evidence supports hypotheses that it makes more credible. His triviality result leads him towards adequacy conditions for the former idea, but his theory of instance confirmation still carries vestiges of the latter idea.

2.3 Formality and Evidential Support

These counterexamples to Hempel's conditions reveal a more fundamental problem. The most obvious response to them is not to reject, e.g., instance confirmation or the Consequence Condition, etc., but to note that they hold in some cases and not in others. The grasshopper counterexample to Nicod's condition from §2.1 is illustrative here. In that case, an instance of a generalisation undermines it, because of some particular features of the instance together with our background knowledge. Likewise in Missing Bushwalker; our background knowledge about the possible locations of the Bushwalker allows information consistent with a hypothesis to undermine it, whereas in a different example information with the same logical relation to the hypotheses under consideration would not have been undermining. (For example, if we initially had reason to think the bushwalker was south of the river, the information might have simply confirmed our original view and left our reasons in favour the various hypotheses unchanged.

This unavoidable role for background knowledge shows that inductive logic cannot be exactly like a logic, for logic is formal while there is no possibility of a fully formal theory of evidential support. Certainly some aspects of evidential support derive from logical features. Every theory I know of respects the Entailment Condition, that entailment is a limiting case of evidential support. But examples in which the logical relations between evidence and hypothesis are constant and yet the evidential support relations vary are multiple. Suppose we are drawing from an urn of known constitution, containing two red and two black balls. Against a background containing information that the draws are made with replacement, the evidence that the first two draws were red followed by red provides no support for the hypothesis that the next draw will be red. Against a background containing the information that the draws were made without replacement, that same evidence is conclusive support for the hypothesis that the next draw is black.

A good account of evidential support will incorporate this explicitly:

> Hempel's account is concerned with a two-place relation ('E confirms H') rather than with a three-place relation ('E confirms H relative to K').... one of the morals ... to draw is that background knowledge can make a crucial difference to confirmation. (Earman 1992: 67)

Some have argued that the non-formality of induction means there can be no general theory of evidential support at all:

> We have been misled, I believe, by the model of deductive logic into seeking an account of induction based on universal schemas. .... Instead inductive inferences will be seen as deriving their license from facts. These facts are the material of the inductions..... Particular facts in each domain license the inductive inferences admissible in that domain—hence the slogan: "All induction is local." (Norton 2003: 648)

In the end Norton's objection seems to boil down to a claim that inductive inference is material, so there can be no topic-neutral theory of what we might infer from a given piece of evidence – not even one that incorporates background knowledge. This however is something that our synchronic conception of inductive logic has already conceded (§1.4). Norton's more controversial point appears to be a kind of rational particularism: that there are no universal principles governing rationally permissible states of belief (Norton 2011: 399–401). For example, he suggests that there cannot be a general theory of how background knowledge mediates between evidence and hypothesis, because the background knowledge itself dictates how background knowledge bears on evidential support.[2] The only response to this line of argument is to propose a general model of evidential support, and to analyse the cases the particularist appeals to in such a way as to show that in fact they are coherent with the general model. That, in any case, is the approach I will adopt.

2.4 Hume's Problem of Induction

Examples showing the 'materiality' of induction are also illustrations of a venerable challenge to the rationality of induction, originally due to Hume (1777). Hume's focus is inductive inference, and his principal concern is normative, not descriptive. What, he asks, justifies our inductive inferences? Perhaps more importantly, under what circumstances does an inductively generated belief amount to knowledge? Hume expresses 'curiosity' about whether inductive knowledge is possible.

> Here then is our natural state of ignorance with regard to the powers and influence of all objects. How is this remedied by experience? It only shews us a number of uniform effects, resulting from certain objects, and teaches us, that those particular objects, at that particular time, were endowed with such powers and forces. When a new object, endowed with similar sensible qualities, is produced, we expect similar powers and forces, and look for a like effect. From a body of like colour and consistence with bread, we expect like nourishment and support. But this surely is a step or progress of the mind, which wants to be explained. When a man says, *I have found, in all past instances, such sensible*

[2] To foreshadow, he is arguing here against the Bayesian approach Section 3 where the representation of belief makes use of probabilities, and where any rational state of belief corresponds to a probability function; he thinks background knowledge can tell us not to represent a given situation probabilistically.

> *qualities conjoined with such secret powers*: And when he says, *similar sensible qualities will always be conjoined with similar secret powers*; he is not guilty of a tautology, nor are these propositions in any respect the same. You say that the one proposition is an inference from the other. But you must confess that the inference is not intuitive; neither is it demonstrative: Of what nature is it then? To say it is experimental, is begging the question. For all inferences from experience suppose, as their foundation, that the future will resemble the past, and that similar powers will be conjoined with similar sensible qualities. If there be any suspicion, that the course of nature may change, and that the past may be no rule for the future, all experience becomes useless, and can give rise to no inference or conclusion. It is impossible, therefore, that any arguments from experience can prove this resemblance of the past to the future; since all these arguments are founded on the supposition of that resemblance. Let the course of things be allowed hitherto ever so regular; that alone, without some new argument or inference, proves not, that, for the future, it will continue so. ... My practice, you say, refutes my doubts. But you mistake the purport of my question. As an agent, I am quite satisfied in the point; but as a philosopher, who has some share of curiosity, I will not say scepticism, I want to learn the foundation of this inference. (Hume 1777: ¶4.21)

Hume does not frame his objection as one concerning evidential support. But it is not difficult to see that he is here raising a challenge to the foundations of inductive logic. Our conception of inductive logic suggests that E's supporting H is a necessary condition on H coming to be known inductively (i.e., via an inference prompted by the acquisition of knowledge that E). This cannot plausibly be a sufficient condition: consider lottery cases, in which some evidence concerning the workings of a large fair lottery can overwhelmingly support the hypothesis that a single ticket without, intuitively, amounting to knowledge (Hawthorne 2004). But if evidential support is to play any role in inductive inference, it will surely be necessary that one can acquire inductive knowledge only by inference from supporting evidence.

As Hume would immediately point out, however, such an inference rests on the supposition that E supports H – in his case of interest, on the supposition that evidence about the past 'is a rule' for hypotheses about the future. For the hypothesis to be known as a result of this inference, we need to know that supposition to be true. But in many cases, Hume argues, we do not. When Hume emphasises the possibility 'that the course of nature may change', he is not merely pointing out a difference between induction and deduction. He is pointing out that our grasp on relations of evidential support is itself dependent on our evidential history. More concisely: any knowledge of relations of evidential support presupposes prior knowledge of how past evidence supports hypotheses about evidential support (Howson 2000: 10–15).

One response to this argument might be to challenge Hume's assumption that knowledge of evidential support is contingent on one's evidential history. Perhaps relations of evidential support are after all 'demonstrative'? This seems not to be the case. Consider a slight modification of one of Good's cases showing the falsehood of Nicod's Condition:

> Suppose that we know we are in one or other of two worlds, and the hypothesis, H, under consideration is that all the crows in our world are black. We know in advance that in one world there are a hundred black crows, no crows that are not black, and a [b]illion other birds; and that in the other world there are a thousand black crows, one white one, and a million other birds. A bird is selected equiprobably at random from all the birds in our world. It turns out to be a black crow. This is strong evidence … that we are in the second world, wherein not all crows are black. Thus the observation of a black crow, in the circumstances described, undermines the hypothesis that all the crows in our world are black. (Good 1967: 322)

This case shows that the evidence – seeing a black crow – is variably supportive of the hypothesis that all crows are black, depending on the background conditions. Whatever evidence we have that suggests we are in one or the other of these two worlds, is thereby evidence for different hypotheses about the evidential significance of the black crow. Seeing a decent number of birds might be strong evidence that we are in the first world, rather than the second. And so if our past history contains substantial bird experience, then we might come to accept that seeing a black crow supports the hypothesis that all crows are black. On the other hand, a different course of experience, less rich in avian content, might support the second hypothesis about which world we are in, and a second hypothesis about what the observation of a black crow supports.

Our modification of Good's case illustrates two key features of inductive inference. One we have encountered already in §2.3, that evidential support cannot be understood as an internal relation between evidence and hypothesis, but is mediated by background conditions (whether we are in world one or world two). The other is that our evidence for background conditions might itself be unrepresentative (perhaps, by chance, we see a significant number of birds, suggesting falsely we are in world one), and that ought to call the reliability of our inductive inferences into question for us (Kelly 1996: 46ff).

Hume's resolution of this problem is characteristically bold. He denies that 'inferences from experience' are examples of reasoning at all (Hume 1777: ¶5.4).

> All belief of matter of fact or real existence is derived merely from some object, present to the memory or senses, and a customary conjunction between that and some other object. … This belief is the necessary result of placing the mind in such circumstances. It is an operation of the soul, when we are so situated, as unavoidable as to feel the passion of love, when we receive benefits; or hatred, when we meet with injuries. All these operations are a species of natural instincts, which no reasoning or process of the thought and understanding is able, either to produce, or to prevent. (Hume 1777: ¶5.8)

Hume thinks we do have a fairly widely shared conception of the support provided by specific pieces of evidence. But he thinks these depend on a habit or custom we have acquired as an effect of experience, rather than being the product of reasoning, and accordingly they are not subject to rational evaluation (Hume 1777: ¶5.5). We have managed to form persistent beliefs about evidential support, which contribute to inductive reasoning, providing a 'foundation' for this type of inference without necessarily justifying it.

What does Hume's naturalistic explanation of our inductive habits tell us about evidential support? In one way, it merely underscores the result from §2.3 that evidential support must be understood as a three-place relation, but clarifies that the third relatum is not merely 'background knowledge', but background knowledge that suffices to yield a standard of evidential support. The fact that evidence bears on hypotheses differently in different circumstances suggests that a full understanding of evidential support will involve an explicit relativisation to some sort of background model of evidential support. Hume's suggestion that inductive inference in the narrow sense requires the prior assumption 'that the future will resemble the past, and that similar powers will be conjoined with similar sensible qualities' (1777: ¶4.21) is an example of what such a model might be. This principle tells you to treat past evidence as a representative sample. A principle that told us to treat past evidence more cautiously, as a potentially unrepresentative sample, would license quite different claims of evidential support.

We don't need to follow Hume in his inductive scepticism, or his preference to explain rather than justify our inferential tendencies. It might well be that we can, ultimately, find a justification for accepting a particular conception of evidential support. For example, perhaps some conception of how evidence bears on hypotheses is justified by default, or we might have justification for some inductive assumptions 'without being in a position to cite anything that could count as ampliative, non-question-begging evidence for those beliefs' (Pryor 2000: 520, though he is talking about perceptual beliefs in this context). Or perhaps we acquire our inductive habits of thought by testimony or immersion in a scientific community, and justification comes subsequent to the fruits of our epistemic endeavours. Perhaps we are justified simply because 'the world is so constituted that inductive arguments lead on the whole to true opinions' (Ramsey 1926: 93).

However such justification might be acquired, it is fruitful to separate the issue of how evidence bears on hypotheses from the question of the *structure* of evidential support. The latter issue can be pursued to a significant extent by characterising the connections between evidence, hypotheses, and standards of evidential support, without taking a stand on what the actual evidence is, what the actually live hypotheses are, or what the actual standards of evidential support are. Because this investigation concerns the formal features of this three-place relation 'E supports H relative to standards S', it is a deserving bearer of the label 'inductive logic'. (Though to illustrate its significance, it will be helpful to give at least some examples of how the relata can be concretely filled in.)

That is how we will pursue inductive logic in this book. In Section 3, 4 we will look at *Bayesianism*, a particularly influential and fruitful approach to the formal features of evidential support. Bayesians understand a conception of evidential support probabilistically: to regard the total evidence as supporting H is to adopt a probability model according to which H is credible; given that same model, to regard E as confirming H is to regard H as more likely given E than H otherwise. In §4.7, we will return to this issue of justifying inductive assumptions, accompanied by an extensive investigation into whether there are significant prior constraints on what sorts of probability models we may permissibly adopt Section 5; this will include a brief consideration of Goodman's (1954) 'new riddle' of induction (§5.4).

Section 3

Probability and Evidential Support

In this section, we discuss how probability can be used to understand evidential support. After considering the motivating metaphor of the weight of evidence (§3.1), suggesting that the weight of evidence fixes the prospects of a given proposition. Prospects, in turn, are relative to a perspective, what comes in view from the vantage of a given body of evidence; epistemic perspectives are described in §3.2, and their connection with the rationalization of idealized bets is made clear in §3.3. Constraints on the proper evaluation of bets then constrain legitimate measures of prospects (§3.4). Indeed, the constraints ensure that every legitimate measure of prospects is an evidential probability function §3.5; we spend some time characterising the principles governing these probability functions. The argument from §3.4 is a version of a Dutch Book Argument, and I delve into the differences between my version and more standard versions in §3.6, and investigate accuracy-based arguments for evidential probability in §3.7. In §3.8 we turn to conditional probability, the prospects of hypotheses given other claims, which will be central to our conception of evidential support and confirmation. Finally, in §3.9 I turn to the epistemic question of how to choose or rationalize a choice of epistemic perspective, and touch on the topics of epistemic deference and the debate over permissivism in epistemology.

3.1 Weighing Evidence

A theory of evidential support aims to capture a number of widely shared beliefs about what evidence does for us. For example: that unexpected or novel predictions should favour a hypothesis more than banal or familiar predictions; or that simpler hypotheses, other things being equal, are favoured by the evidence. Behind these beliefs seems to be a persistent metaphor: evidence as a mass stuff, something that can be gathered and, once collected, weigh in favour of a hypothesis and against its rivals.

This metaphor of the 'weight of evidence' is presumably carried over from the role given to evidence in the legal context, where evidence favouring and disfavouring a hypothesis about guilt is weighed by the scales of justice. This model in turn suggests a natural quantitative approach to evidential support as a measure of the weight of evidence. And since the common standard applied in criminal trials is that the evidence must render the hypothesis of guilt beyond reasonable doubt before a juror should vote to convict, evidence must therefore be weightier the more it eliminates doubt or uncertainty.

This metaphor applies even in advance of the evidence being gathered. A striking new piece of evidence is often recognised as such because its potential was known prior to it coming in to evidence. A measure of evidential support indicates the prospects of hypotheses, in that a hypothesis which is strongly supported by some potential evidence has a greater prospect of turning out to be true if that potential evidence is gathered.

The picture of evidential support as measuring how potential evidence for a hypothesis would, hypothetically, reduce doubts about it, combined with these widely held beliefs about the role of evidence, will jointly constrain the project of inductive logic. Our core analysis will be motivated by the connection with doubt and certainty; the merits of that analysis are revealed through its success in explaining the truth (or apparent truth) of those commonplace beliefs.

3.2 Prospects and Perspectives

Evidential support measures the extent to which some potential evidence improves the prospects of some hypothesis. The prospects for a hypothesis are

– again metaphorically – what comes into view in light of some body of evidence. What is visually in prospect for you depends on where you are, on your perspective. Likewise, what is *epistemically* in prospect depends on a particular perspective on the space of possibilities. Camp (2019) characterizes occupying an epistemic perspective as having 'an open-ended disposition to characterize: to encounter, interpret, and respond to some parts of the world in certain ways' [p. 24]. That is what it is to *occupy* a perspective. A perspective itself need not be embodied in anyone's dispositions – an epistemic perspective exists independently of any agent occupying it – but it must contain sufficient information to characterize such dispositions. To fulfil its role, I understand an 'epistemic perspective' to involve at least the following components:[1]

1. A representation of the space of possibilities;
2. A representation of the total evidence currently on hand, and
3. Some policy, or set of standards (Schoenfield 2012: 199), that capture, numerically, the significance of potential evidence for various entertainable hypotheses.

Let us expand on these components. 'The space of possibilities' captures all the ways that things could turn out to be, according to a given epistemic perspective. It is natural to take this space of possibilities to have the structure of a *field* of propositions (Eagle 2011: 1–3; Hájek and Hitchcock 2017: 17). That is, for each elementary outcome that could occur, according to that perspective, there is a proposition in the field to the effect that the outcome comes to pass; there is a trivial outcome, represented by a trivially true proposition (e.g., that something or other comes to pass); and whenever the field of propositions includes P and Q, it also includes $\neg P$, $(P \vee Q)$, and $(P \wedge Q)$. A maximally specific possibility will be represented by a logically complex proposition that fixes, for each possible outcome, whether or not it occurs. This will play the role of a 'possible world' from that perspective. It must be acknowledged that a different perspective might discriminate outcomes that a given perspective treats as indistinguishable, and thus different perspectives needn't agree on how to analyse any given proposition into more specific out-

1 This is not how Camp pursues the idea.

comes.[2] So there is no guarantee that all perspectives will agree on what the elementary possibilities are, or on which propositions are possible. We will not in what follows focus on these possible differences between perspectives, and will assume that all the perspectives we consider share the same space of possibilities, and are capable of representing any hypotheses or piece of evidence we might need to consider. The question of how to understand cases where an agent expands their conception of what is possible is philosophically very rich and much discussed (sometimes under the label of 'partition sensitivity'), but lies beyond our scope (Paul 2014; Pettigrew 2020a).

We will also set aside for the most part interesting questions about whether propositions are to be understood as only as fine-grained as possible worlds, in which case every metaphysically necessary truth will be true throughout the space of possibilities (Stalnaker 1984: 2). The alternative is to adopt any of a number of conceptions of propositions on which there are epistemically possible propositions corresponding to 'impossible worlds' (King, Soames, and Speaks 2014, esp. pp. 33-44). (For example, if we thought there was some epistemic possibility of a unknown mathematical truth being false, we'd need to be able to entertain an impossible world in which that proposition were true.) I hope to avoid taking a stand on this question of the grain of propositions, but as I will assume that the objects of any epistemic perspective are propositions, our discussion is somewhat hostage to the results of this ongoing discussion in the philosophy of language.

An epistemic perspective must also represent a body of total evidence. Evidence constrains which possibilities in the space of possibilities are left open and which excluded. Because evidence is propositional, evidence narrows down a location in the space of possibilities by excluding those possibilities in which propositions inconsistent with the evidence are true (Stalnaker 1984: 120). So an epistemic perspective can represent the current total evidence by indicating in some way a region of the space of possibilities which, from that perspective, might be actual – those consistent with the evidence.

Finally, and most importantly for inductive logic, an epistemic perspective must represent in some way the bearing of the evidence on the remaining

2 For example, the propositions 'Heads' and 'Tails' might represent the elementary possibilities according to some very limited perspective that is unable to distinguish between worlds in which a specific coin lands the same way. The field of propositions given these outcomes will be the set 'Heads', 'Tails', 'Heads or Tails', 'Neither Heads nor Tails'.

live hypotheses. It must discriminate between those hypotheses consistent with the evidence – those in the region of open possibilities. We will assume that this discrimination is effected numerically; so an epistemic perspective assigns numbers to hypotheses that somehow reflect their support by the evidence. As Horwich puts it, 'our inductive practice may be represented by a function which specifies, for any evidential circumstance, the permissible degrees of belief in any statement' (1982: 79) – such a function will be an epistemic perspective, noting that it is not to be identified with one's attitudes, but circumscribes permissible attitudes.

In so doing, an epistemic perspective thus represents a possible way of 'proportion[ing] belief to the evidence' (Hume 1777: ¶10.4). A proportioning will in general assign different numbers to hypotheses, even those it regards as on balance supported by the evidence. That evidential support varies in extent is fairly commonsensical. Suppose that my current evidence supports the claim that I will be heading home from the office shortly, and that it also supports the claim that the weather next weekend will be very hot. It would be foolish to deny that it may support the former to a considerably greater extent than it supports the latter. That this variability in evidential support can be quantified numerically is a reasonable starting point.

Someone's actual attitudes may not correspond to any epistemic perspective; for example, they might not have any coherent attitude to a hypothesis that would enable a single number to be assigned as to which the evidence supports it; perhaps at best an imprecise range of attitudes could be ascribed to them (Jeffrey 1983: 139–40). And an epistemic perspective may or may not correspond to anyone's actual doxastic attitude. But an epistemic perspective does appear to be something a rational agent could have as an ideal for a coherent doxastic attitude. To put it in explicitly normative terms, an epistemic perspective comprises a structure that reasonable doxastic attitudes must approximate (perhaps each reasonable indeterminate attitude can be precisified into a determinate epistemic perspective), and provides a regulative ideal for such attitudes (see §3.9). So we will be assuming in a sense that all epistemic perspectives are rational, and trying then to establish the properties that distinguish epistemic perspectives from other purported evaluations of the prospects of hypotheses. Epistemic perspectives rationalize the doxastic dispositions of the agents who occupy them.

3.3 Prospects and Bets

If an epistemic perspective represents the prospects of hypotheses, it must be subject to certain norms about the evaluation of prospects. One way to bring out these norms is to consider the ideal evaluation of *bets* (Ramsey 1926; de Finetti 1937; Pettigrew 2020b).

Real-life gambling behaviour is psychologically complicated and morally fraught. The declining marginal utility of money (Pettigrew 2020b: 17–19), the difficulty of finding enough parties and counterparties to every possible bet, the fact that some outcomes will be resolved only after the agent's lifespan, the fact of risk-averse and risk-seeking individuals (Kahneman and Tversky 1979; Buchak 2013), all make it difficult to draw direct conclusions about the prospects of P from an agent's (un)willingness to bet on P. Nevertheless, behind any reasoned decision to bet lies an epistemic perspective: some evaluation of the prospects of the propositions one is betting on, given the evidence one has. That the prospects of a proposition are decent can go some way towards *justifying* an agent's a decision to bet, without constituting, or being constituted by, that willingness to bet.[3]

If an agent's epistemic perspective is to justify their dispositions to bet, what conditions must that epistemic perspective satisfy? The evaluation of a given bet as attractive or otherwise results from combing an epistemic perspective with an exogenously given assignment of values to outcomes. If there are constraints on the acceptability of evaluations, those constraints may reveal structural constraints on any epistemic perspective that is properly judging the prospects of outcomes.

A *bet* on a proposition H is a right to receive s units of value if H turns out to be true, and nothing otherwise. The price of the bet is the value x that is assigned to that right. For each bet, there is a *counterbet*, which is a bet against H that yields s units if H turns out to be false and which costs $s - x$ units. The payout of s utiles is the total staked, the sum of the units of value contributed jointly by the bet and counterbet, $s = x + (s - x)$. For simplicity, let us restrict attention to bets where the total stake s is 1 unit. We can treat a counterbet against H at $s - x$ as equivalent to a bet on $\neg H$ at that same price.

[3] There are many epistemic perspectives, so whether an agent is justified in betting will also depend on whether they were epistemically justified in adopting or deferring to a given epistemic perspective.

This is justified by this basic principle of *equivalence*: any bets which have the same payoffs in the same circumstances must rationally be assigned the same price. A bet against H that costs $s - x$ to yield s pays off when H is false; as does a bet on $\neg H$ with the same cost and yield.[4]

Real agents only enter into bets they regard as favourable. A bet on H is favourable, intuitively, if the prospect of gaining the payout s more than compensates for the risk of losing what you have staked x. We have assumed that an epistemic perspective represents prospects numerically. Suppose that relative to a given perspective the prospect of H is p. A bet is favourable if the payout, weighted by the prospect of getting the payout, exceed the cost: $ps > x$. A bet is unfavourable if $ps < x$; a case where the prospect of gain is outweighed by the risk of loss. An agent is epistemically justified in entering into a bet only if the bet is favourable-from-their-perspective. (This is a necessary condition for all-things-considered justification for a bet, but is not sufficient.)

What if the prospect of gain and the risk of loss are exactly balanced according to some perspective? Such a bet is *neutrally* priced.[5] The potential bettor is not justified in entering into a neutrally priced bet; the prospect of gain is too little. But nor is the bet unfavourable; the prospect of loss isn't sufficiently great to motivate a favourable evaluation of the counterbet. So a neutrally priced bet is one that, given the prospects of the bet, it would not be reasonable to prefer to accept the bet or its counterbet over the *status quo*.[6] *Thus a neutrally-priced bet will represent an accurate evaluation of the prospects of H according to an epistemic perspective*: it is a bet on H that is calibrated to the degree to which the evidence supports H.

To those tempted by some sort of instrumentalism, it might be appealing to try and assign prospects to hypotheses, according to an epistemic perspective, by trying to behaviourally elicit a neutral betting price. This is unlikely to

4 This assumes that the logic of negation, and structure of the space of propositions, is classical; constructing probabilistic epistemic perspectives given a non-classical logic involves several diverting challenges (Williams 2016).

5 A neutrally priced bet is often called a 'fair' bet, but this has an unhelpful contrast with 'unfair'; not all non-neutral bets are unfair to the bettor.

6 Of course, an agent who is risk-seeking may wish to take on such a bet anyway, but such a course of action isn't epistemically justified by the prospects of the proposition bet on, though it may all-things-considered benefit the agent if they sufficiently value the thrill of gambling.

succeed, given all the confounding factors in actual betting behaviour. (This is despite Ramsey's suggestion that a neutral price reflect those odds someone would 'just take' (Ramsey 1926: 72), or Kyburg's insistence that 'how seriously someone believes what he says he believes' is elicited by inviting him 'to put his money where his mouth is' (1983: 64).) Indeed it is complicated to deduce even an agent's own mental state from the betting prices they offer, in general (Bradley and Leitgeb 2006), let alone the epistemic perspective legitimising their attitudes. Hence I will suppose that there are epistemic perspectives, and that they serve in some complex way to justify dispositions to accept particular bets (Howson and Urbach 1993: 76–77).

3.4 Evaluating Prospects

To assign a neutral price to a bet would involve fixing on some particular epistemic persepctive. The general structural principles on epistemic perspectives we are concerned with can be uncovered without making assumptions about neutral prices except that they exist. We uncover them by noting that unless certain constraints are placed on the numerical evaluation of prospects, certain bets (and packages of bets) seem to yield an incoherent evaluation of their value. Presented one way, the bets have a certain neutral price; presented another way, they are assigned another non-neutral price. So a single neutral price doesn't exist unless we constrain the acceptable numerical representation of evidential support in some way.

For example, suppose an epistemic perspective assigned a degree of support $-\delta$ to a proposition H, where $\delta > 0$. Such an perspective entails that the neutral price for a bet *on* H that pays 1 unit as $-\delta$. But at that price, there is an advantage to purchasing the bet – the price is negative, so, win or lose, represents an advantage over the *status quo*. So the price is not after all neutral. A perspective that assigns a negative number as the degree to which the evidence supports H leads to a situation where the agent cannot assign a coherent neutral price, because their assessment of the prospects of H leads them to be indifferent to accepting this bet, but an assessment of the possible payoff of the bet should lead them to be positively disposed to accept it. So this epistemic perspective turns out to be unable to assign a single value to the bet. To assign a neutral price to a bet, no perspective could assign a degree of support less than zero. That is, every epistemic perspective must satisfy:

Non-negativity The degree of support provided by any body of evidence for a hypothesis H must be greater than or equal to zero.

This argument could be resisted if the neutral price of a bet depends on how the bet is described. If bets are individuated very finely, it might be that there is an advantage over the *status quo* to the bet described one way, and not another. But in fact it is not possible to see how this would work. The advantage or disadvantage of a bet is due to the conditional rights to valuable goods they convey. If one bet conveys at least as much as another, no matter what, one cannot rationally prefer the second to the first, regardless how it is described. To do so would be 'absurd', as Ramsey says: any account of prospects 'which broke [this principle] would be inconsistent in the sense that it violated the laws of preference between options, such as that preferability is a transitive asymmetrical relation, and that if α is preferable to β, β for certain cannot be preferable to α if p, β if not-p' (Ramsey 1926: 78). Since in this case $1 + \delta$ is preferable to δ, so the bet on H that costs $-\delta$ is preferable to δ for sure; and δ for sure is preferable to the status quo, transitivity of preference ensures that the bet cannot be neutrally priced at $-\delta$.

As this case shows, an epistemic perspective violating Non-negativity justifies inconsistent evaluations of the very same option. That inconsistency is what excludes it from being a genuine perspective on the prospects of H; the betting setup is simply a way to make this inconsistency manifest. The same sort of argument can be given for other constraints on epistemic perspectives.

Suppose that an epistemic perspective assigned a degree of support $\delta < 1$ to a trivial hypothesis H – one that was necessary, according to that perspective (obtaining in every possibility considered by that perspective). Such an epistemic perspective would regard as neutrally priced a bet on H priced at δ. Such a bet is certain to pay off; for such a perspective regards H as already true. So this bet cannot be neutrally priced at δ; for there is a guaranteed advantage to purchasing it, hence it is not to be regarded with indifference from the *status quo*. So for the same reasons as above – namely, that coherent epistemic perspectives should assign a single neutral price to a bet – no epistemic perspective could assign a degree of support to a certain outcome that was less than 1:

Normality The degree of support provided by any body of evidence for a trivial hypothesis (necessary in light of that evidence) is 1.

Suppose that there are two hypotheses, H and H', such that the truth of either one excludes the truth of the other, according to some epistemic perspective. Suppose an epistemic perspective assigns a degree of support α to Q, and degree of support β to R, and a degree of support δ to their disjunction $H \vee H'$, but where $\delta \neq \alpha + \beta$. The neutral prices for bets on these propositions H, H' and $H \vee H'$ are fixed by those degrees of support. Consider the book of bets consisting of bets on H and H' that each pay 1 unit. This book of bets the payoffs depicted in Table table 3.1. It is easy to see from the table that where $H \vee H'$ is true (the top two possibilities), the payoff is $1 - (\alpha + \beta)$, and when $H \vee H'$ is false, the payoff is $-(\alpha + \beta)$. But the neutral price for a bet on $H \vee H'$ is δ, leading to a different payoff of $1 - \delta$ when true, and $-\delta$ when false. So the book of bets pays out the same amount, in the same circumstances, as the individual bet on the disjunction – but has a different neutral price. Clearly this either violates the equivalence requirement to assign the same neutral price to bets that pay off in the same circumstances, or violates the requirement on epistemic perspectives that we assign a single neutral price to a given bet. Hence:

Table 3.1: Payoffs for the book of bets on H and H'.

H	H'	Bet on H	Bet on H'	Total payoff	$H \vee H'$
T	F	1-α	$-\beta$	$1-(\alpha+\beta)$	T
F	T	$-\alpha$	$1-\beta$	$1-(\alpha+\beta)$	T
F	F	$-\alpha$	$-\beta$	$-(\alpha+\beta)$	F

Additivity The degree of support provided by any body of evidence for a disjunction of mutually exclusive hypotheses is the sum of the degrees of support provided to each hypothesis individually.

One further principle is implicated in establishing Additivity as a requirement on epistemic perspectives: the *package principle* that the neutral price for the book is the sum of the neutral prices for the bets, i.e., 'that the value I set on them together is the sum of the values I set on them singly' (Schick

1986: 113; Earman 1992: 42).[7]

3.5 Probabilities and Degrees of Support

The argument of the previous section reached the conclusion that coherent epistemic perspectives, those justifying a single neutral price for a bet on each proposition in their scope, must satisfy Non-Negativity, Normality, and Additivity.[8] Recalling that an epistemic perspective is a numerical function from a field of propositions (§3.2), this establishes that epistemic perspectives must be — mathematically speaking — probabilities. For satisfying these constraints is sufficient for a function to be a probability function (Kolmogorov 1933; Eagle 2011: 1–4; Hájek and Hitchcock 2017).[9]

In light of this, degree of support can be understood as the extent to which a given body of evidence makes a hypothesis likely to be true. Strongly supported propositions are probable in light of the evidence; weakly supported propositions are improbable in light of the evidence. Thus we may conclude that each epistemic perspective also constitutes an *evidential probability* function, assigning numbers to hypotheses, representing degrees of support, in light of the total evidence that perspective is committed to. In line with §3.2, the total evidence according to an evidential probability function comprises those propositions which are assigned probability 1. This usage might be slightly revisionary, in that it will include logical consequences of evidence as evidence, and logical truths and other trivial propositions will also be evidence. No harm comes from liberalising our conception of evidence from the usual idea that it only includes propositions learned from experience, to the idea that it

7 This package principle is plausible when one evaluates a package of bets offered all at once, especially given that we are considering the price that is justified, not the price that someone's wallet can withstand. (Add sufficiently many advantageously priced bets together and there will come an advantageously priced bet that one would purchase offered alone, but cannot afford.) Yet as Schick points out it is far less obvious when bets are offered sequentially (1986: 116–18). The fact that one has accepted a given bet is new information, and betting that would have been justified from one's previous epistemic perspective may not continue to be justified from an updated perspective.

8 In fact we need an additional premise here, that if you do obey these principles, you can assign a single neutral price: that premise is true (Kemeny 1955; Lehman 1955).

9 I will not review the mathematics of probability in this brief volume, but I will try to be explicit about any results proved below, and there are many accessible and brief presentations of the key results on which I will rely (Hacking 2001: 23–78; Eagle 2011: 1–24; Hájek and Hitchcock 2017).

includes at a given point any proposition on which we rely in evaluating other claims with which we are confronted. 'Taking the evidence into consideration' doesn't mean only taking things you happen to have learned empirically into consideration.

Our approach to evidential probability has precedent in the literature. As we argued in §3.2, epistemic perspectives rationalize the dispositions of their occupants to judge hypotheses. Thus we agree with Williamson (2000), who suggests that evidential probability represents the degree to which 'the evidence tells for or against the hypothesis' (2000: 209), rather than reflecting anyone's actual credences.[10] Climenhaga puts it well:

> the distinctive claims of the degree-of-support interpretation [of evidential probability] are that probabilities are mind-independent relations between propositions and that probabilities constrain rational degrees of belief. (Climenhaga 2023: 3)

The tradition of *objective Bayesianism* that Climenhaga and Williamson represent goes back at least to Keynes (1921), Johnson (1932), and Carnap (1962). It has however almost invariably been accompanied with an additional and not necessarily welcome commitment to there being such a thing as '*the* degree to which evidence supports a hypothesis'. While each epistemic perspective on my view articulates a conception of how evidence bears on hypotheses, there is no commitment to the existence of a 'best' perspective. Perspectives – for all we've established so far – might disagree with one another on the space of possibilities, on the background evidence they incorporate, or on the significance they attach to it. We return to the question of whether there is a unique best perspective in §3.9 below and in Section 5. For now, I would wish to distance the view defended here both from the 'subjective' Bayesian, who identifies epistemic perspectives with the credences on individual agents, and from the objective Bayesian who accepts a unique best perspective. The view I'm defending could be called *non-subjective Bayesianism*.

10 Williamson himself has a more restrictive view, that evidential probability should reflect 'something like the intrinsic plausibility of hypotheses prior to investigation' (Williamson 2000: 211); we accept that among the epistemic perspectives are those one can come to occupy *after* some investigation.

It is clear that among the possible evidential perspectives are many that assign a zero probability to contingent propositions. Indeed, any piece of contingent evidence that has already been taken into account, forming part of the total evidence according to some epistemic perspective, will be assigned probability one by that perspective. Accordingly, we do not require that every epistemic perspective Pr satisfy the principle of *Regularity* (or 'strict coherence'), that the only propositions H for which $\Pr(H) = 0$ are those that are necessary falsehoods, true at no possible world (Lewis 1986a: 88).

We do not even require the weaker principle that any epistemically possible proposition should be supported to some positive extent by the evidence. For sometimes there are too many possibilities for them all to be assigned some positive probability. Consider an infinite sequence of independent tosses of a fair coin. (Suppose an epistemic perspective according to which such things are possible, and suppose moreover that the tosses happen increasingly swiftly so the whole infinite sequence is completed in a finite time.) Let T be the proposition that every toss lands 'tails', and let T_n be the proposition that the first n tosses land tails. By fairness and independence, $\Pr(T_n) = \frac{1}{2^n}$.[11] If $\Pr(T) > 0$, then there will be some k such that $\Pr(T_k) = \frac{1}{2^k} < \Pr(T)$. But clearly T entails every T_n, and hence $\Pr(T) \leq \Pr(T_k)$.[12] Contradiction; so $\Pr(T) = 0$, despite being an epistemically possible outcome in the envisaged scenario. Some respond here that we have too narrow a conception of the available numbers. Had we appealed to 'infinitesimal' numbers, greater than zero but less than any real number, we could allow the probability of T to thread the gap between 0 and the probability of any initial subsequence of T (Lewis 1986a: 88). Infinitesimals might help in this case; Williamson (2007: 175–76) argues that they don't help in general. Consider a *final* subsequence of T, an infinite sequence T_{-k} omitting finitely many (k) initial elements of T. Such a sequence T_{-k} must have the same probability as T, because it is structurally indistinguishable from T. But independence entails that $\Pr(T) = \frac{1}{2^k} \Pr(T_{-k})$, which entails that $\Pr(T) = 0$. So it seems mathemat-

11 This appeals to the *multiplication theorem* that for independent events A and B, $\Pr(A \wedge B) = \Pr(A)\Pr(B)$; independence is clarified below (§3.8).

12 This follows from the *consequence theorem* that the probability of a logical consequence of H is no less than the probability of H.

ically unavoidable that we assign some contingent hypothesis a zero degree of support.

Once that is recognised, there seems no grounds for insisting that open possibility has to be assigned some positive degree of support. Epistemic perspectives will be notably diverse in their opinions concerning the degree to which various propositions are supported. There is no reason to insist that nevertheless they should all agree that epistemic possibilities should be supported to some extent. This means that we will permit pairs of epistemic possibilities that are *incommensurable*: that do not agree antecedently in which are the live hypotheses that might turn out to be true, that have some positive prospect. We will return to this issue later when we consider 'convergence of opinion' theorems briefly in Section 5.

I argued that every epistemic perspective respects Additivity. Many approaches to mathematical probability adopt a stronger axiom:

Countable Additivity Where $H_1, H_2, \ldots$ is some denumerable collection of mutually exclusive propositions, and $\bigvee_i H_i$ is a proposition true iff exactly one H_i is true, then $\Pr(\bigvee_i H_i) = \sum_i \Pr(H_i)$.

Countable additivity also requires that there is always such a proposition among those an epistemic perspective is defined over, which requires that the set of propositions themselves be closed under denumerable disjunction.

There are certainly some epistemic perspectives which respect countable additivity. A countably additive function is *a fortiori* also additive. But should we impose it on every epistemic perspective? Many have found the following consequence of Countable Additivity problematic:

No Uniform Support If $H_1, H_2, \ldots$ partition the space of possibilities (being mutually exclusive and jointly exhaustive, so that $\bigvee_i H_i$ is true in every possibility), then there is no uniform assignment of degree of support to each H_i.[13]

[13] Because $\bigvee_i H_i$ is necessary, it is supported to degree 1, by Normality. By countable additivity, $\sum_i \Pr(H_i) = 1$. If there were some ϵ such that $\Pr(H_i) = \epsilon$ for each H_i, then $\epsilon > 0$, because adding zero countably many times results in zero. But if $\epsilon > 0$, then there is some j such that $\Pr(H_1) + \cdots + \Pr(H_j) > 1$. By additivity, $\Pr(H_1 \vee \cdots \vee H_j) > 1$, and yet $\bigvee_i H_i$ is a consequence of $(H_1 \vee \cdots \vee H_j)$.

This result shows evidence cannot equally support a denumerable collection of hypotheses. Suppose we are told nothing more about some rational r than that it is of the form $1/n$. You might think that evidence supports every hypothesis about the identity of r equally; but the above result excludes that. There are ways of assigning degrees of support to a denumerable partition, but they are non-uniform.[14] De Finetti objected that a formal axiom ought not to require a non-uniform distribution of support, but ought to be 'only imposing formal conditions of coherence' (de Finetti 1974: 122; see also McGee 1999; Bartha 2004).

De Finetti suggests that, to secure uniformity, we ought to assign zero degree of support to each element of a denumerable partition (see also Kadane and O'Hagan 1995). In the absence of countable additivity, this does not determine the probability of their infinite disjunction. We are free to set the degree of support of that disjunction independently; it clearly deserves maximal degree of support since it is true in every possibility. This assignment of degrees of support is incoherent in the presence of Countable Additivity. So if it seems coherent on reflection that we could have uniform support for each hypothesis about r, so much the worse for Countable Additivity .

There does exist a betting argument for Countable Additivity (Williamson 1999; Pettigrew 2020b: §2.6), similar to those offered in [@ sec:prospects]. This betting argument requires a way to calculate a neutral price for a countable collection of bets on hypotheses H_i, and a price for a bet on an infinitary disjunction. Thus it requires a generalisation of the package principle which states that the neutral price of a countable collection of bets is the sum of their individual prices. But this generalized package principle appears subject to counterexample. Consider this case:

Satan's Apple 'Satan has cut a delicious apple into infinitely many pieces, labelled by the natural numbers. Eve may take whichever pieces she chooses. If she takes merely finitely many of the pieces, then she suffers no penalty. But if she takes infinitely many of the pieces, then she is expelled from the Garden for her greed. Either way, she gets to eat whatever pieces she has taken.' (Arntzenius, Elga, and Hawthorne 2004: 262–64)

[14] One way is to assign $\Pr(H_j) = 1/(2^j)$; this sequence of values sums to 1 as needed, but is obviously extremely biased towards hypotheses earlier in the sequence.

In this case, each piece of apple has some positive value to Eve; the neutral price she assigns to each of the 'bets' that pay a piece of apple for sure is positive. But the neutral price to taking them all is negative. Quite what Eve should do is unclear (she seems to be rationally required to leave some apple, but how should she decide which pieces to leave?), but it is clear that the neutral price of the package isn't fixed by the prices of the individual bets. That undermines the betting argument for countable additivity as a requirement on all epistemic perspectives (Pettigrew 2020b: 30). Of course that still permits many epistemic perspectives to satisfy countable additivity, so we will make use of it where appropriate.

3.6 Probabilities and the Dutch Book Argument

Section §3.4 presented a version of what is known as the *Dutch Book Argument* for the conclusion that degrees of support must be probabilities (Howson and Urbach 1993: 75–81; Hájek 2008; Eagle 2011: 28–32; Pettigrew 2020b). Many different versions of this argument have been given since the earliest versions due to Ramsey (1926) and de Finetti (1937).

The typical Dutch Book Argument (DBA) looks rather different from the version I have presented, however. The typical version applies to credences (degrees of belief), not to epistemic perspectives (degrees of support). More importantly, rather than arguing that non-probabilistic assignments to numbers to hypotheses fails to reflect the prospects of those hypotheses, the standard DBA argues that non-probabilistic credences are *practically* unreasonable, because they subject those who act on such credences to a sure loss: 'Dutch Book arguments evaluate the rationality of credences by looking at the quality of the choices that they do or should lead us to make' (Pettigrew 2020b: 1). Obviously there is some connection with choices in the argument in §3.4, because an epistemic perspective is supposed to justify a neutral price for a bet, which may partly rationalize choices to bet at non-neutral prices. But the standard DBA relies on a much stronger link with choice, namely, that non-probabilistic credences rationally require the agent to commit to a package of bets that guarantees a sure loss.

These features of the DBA I've defended help defuse some challenges that face more standard versions. The standard version is accused of being unrealistic (because real agents don't have sufficiently many determinate betting preferences for the argument to succeed), or because probabilism is descriptively inaccurate of real agents (Kyburg 1978). Defenders of the standard DBA have tried to fend off the accusation that it requires us to postulate a plethora of bookies wandering around looking to exploit incoherent agents. But nevertheless it is better to sever credences, and the bets they may lead to, from epistemic perspectives, and the credences they justify.

It is by no means obvious that one's credences rationally require any specific course of action. Two agents with the same credences, one risk averse, one risk-seeking, might rationally act quite differently because of their very different attitudes to the frisson of excitement associated with a risky bet. With no direct route from credence to action, a variant of the DBA which doesn't require the evaluation of actions is preferable. Nor is it obvious, once we turn our attention to guaranteed losses and gains, that non-probabilistic credences are so bad. For probabilistic credences require one to avoid the sure gain that one would be able to secure by taking the other side of each of the neutrally-priced bets in §3.4. E.g., you could secure a gain by having credences that neutrally price a bet on a necessary truth $\top$ at 0.8; that bet will pay off which a guaranteed profit of 0.2. This last result is what Hájek calls 'the Good Book Argument' (2005: 142). The Good Book Argument can be avoided if one reformulates the DBA in terms of dispositions to purchase bets regarded as 'fair-or-favourable' (Hájek 2005: 146).

Better still, however, to proceed without bringing in the apparatus of practical action at all, and to focus on the evaluation of prospects that are a mere part of the justification of action. In this I follow other presentations of 'de-pragmatised' DBAs (Skyrms 1984: 22; Earman 1992: 42; Armendt 1993; Howson 2000: 124–34; though see Maher 1997 for a dissenting voice), all of whom emphasise that non-probabilistic credences involve an inconsistent evaluation of the very same options. Nevertheless, even these authors err in focussing on credences. It may well be that there are good reasons for an agent to evaluate options in an inconsistent way, perhaps because of their practical situation. A risk-averse agent whose neutral price for a bet on a fair coin landing heads is 0.49, and likewise for a bet on it landing tails, inconsistently values the *status quo*, neutrally valuing the tautology $H \vee \neg H$ at 0.98. If there is pragmatic encroachment on belief, such grounds for incoherent credence might be widespread (Kim 2017; though see Jackson 2019 for an argument that credence isn't encroachable to the same extent as belief). So I think it preferable to focus on the epistemic perspectives that rationalize credence, rather than on realized credences themselves. (How do epistemic perspectives bear on credences? I say a little about that in §3.9.)

3.7 Accuracy

The DBA is just one argument for the conclusion that epistemic perspectives ought to represents the prospects of hypotheses probabilistically. More recently a flurry of arguments have been offered that aim to give purely epistemic grounds for requiring evaluations of prospects to be probabilistic. In many ways these arguments might mesh closely with the orientation of our discussion, precisely because they focus on how non-probabilistic prospects violate purely epistemic norms. (I argued that non-neutral betting prices pose the purely epistemic difficulty of failing to assign unique prospects to hypotheses, but I concede it can be more difficult to see that clearly if betting has been mentioned.)

Prominent here are *accuracy* based arguments for probabilism (Joyce 1998; Leitgeb and Pettigrew 2010a; b; Pettigrew 2016). The basic idea is that the ultimate yardstick of an epistemic perspective is how accurate it is, how successful it is at representing how things are. On this conception, an *ideal* epistemic perspective is the one that gets everything right about actuality: the only possibilities it regards as having any prospect of truth are those which are in fact true. Almost all epistemic perspectives are non-ideal, and obviously there would be little use for a notion of evidential support if they were: if all and only truths have positive probability, every truth will have probability 1 and be treated as background evidence by the ideal epistemic perspective; nothing is left unsettled. But an epistemic perspective which is not ideal might nevertheless be more or less close to the ideal. The approach shared by all accuracy-based arguments for probabilism rests then on a particular measure of how close an assessment of prospects is to the ideal (typically, the Brier score of the function at each possibility), and a mathematical result establishing that for any non-probabilistic assessment of prospects, there is always a probabilistic epistemic perspective which is more accurate come what may (it is never less close to the ideal at any world, and closer at some world(s)). Given this result, it is then argued that given the value of accuracy, it a non-probabilistic assessment of prospects cannot be a legitimate epistemic perspective, since no matter what there is always a more successful representation of how things might be.[15]

The substantive tasks for a defender of the accuracy-based argument for probabilism involve (i) establishing that the Brier score is the right way to assess the accuracy of assessment of prospects; (ii) establishing that the accuracy of an assessment of prospects is a good measure of its success; (iii) establishing that the a priori existence of a more successful epistemic perspective is grounds for thinking than a given assessment of prospects is unacceptable; and (iv) proving various mathematical results linking these together. I have nothing to say about item (iv), but some comment on the others:

i. *Why one measure rather than another?* Joyce (1998) tried to give necessary conditions on measurements of accuracy which yield the Brier score, but as Maher (2002) noted, these conditions exclude an alternative measure of accuracy (the 'absolute value measure') essentially by fiat. The key contribution of Pettigrew (2016: ch. 4) is to lay down new conditions that yield the Brier score; these nevertheless remain contestable.
ii. *Why accuracy alone?* Epistemologists since Plato have drawn a distinction between true belief and knowledge, and almost all regard knowledge as the more valuable state (Williamson 2000: ch. 1; Pritchard, Turri, and Carter 2022). True belief is accurate; but if this epistemological consensus is right, there are epistemic values that go beyond accuracy. Maybe epistemic perspectives shouldn't care about justification, reasons-responsiveness, open-mindedness, or epistemic virtue, etc. (or should only care about them insofar as they are analysed as requiring accuracy). But if these potential features of an epistemic state could be important, then it is not yet obvious why an assessment of prospects that was very effective at meeting some of these other epistemic goals might not be more successful, epistemically, than some

[15] Here I follow the clear presentation of the dialectic of accuracy-based arguments from probabilism laid out in Pettigrew (2016), pp. 8–9.

more accurate epistemic perspective. For a concrete example: suppose the good knower is one who defers to experts when they are themselves inexpert, and in so doing happen to defer to an expert whose attitudes are non-probabilistic in some inconsequential way. Perhaps the expert is not doing exactly a they ought, but it seems open to us to say that an assessment of prospects that defers to known reliable experts is epistemically unobjectionable.

iii. *Why dominance?* The existence of an probabilistic epistemic perspective that does no worse than a non-probabilistic assessment, no matter how things turn out, is an epistemic version of dominance reasoning from decision theory. Indeed, there is a close connection between accuracy-based arguments from probabilism and so-called 'epistemic decision theory', where epistemic perspectives are the available options and the measure of epistemic value is accuracy. Two comments. First, the reliance on 'epistemologized' decision-theoretic principles seems to leave the accuracy-based approach as implicated in pragmatic commitments as more traditional Dutch Book approaches. (Unless epistemic value is a kind of *value*, why should decision-theoretic principles hold of it?) Second, dominance reasoning goes awry in cases where there isn't independence of acts from outcomes (Bar-Hillel and Margalit 1972). Are there epistemic decision problems with the same structure (Greaves 2013; Pettigrew 2016: 226–29)? (Perhaps an example might be found in the theory-dependence of observation, starting with James' idea that to get evidence for some hypotheses, we need to 'meet the hypotheses half-way' [#james-1910, §10]; crudely put, the more confident you are in H, the more likely you are to find H-supporting evidence.) It's not clear that this ultimately poses any problem for accuracy-based arguments – because probabilism remains an important ingredient of epistemic consequentialism – but some premise in place of dominance seems required.

The upshot of this discussion is that considerations of accuracy provide a compelling further line of argument for probabilism about epistemic perspectives, but one that nevertheless faces challenges not wholly unlike those faced by the argument from unique epistemic evaluation from §3.4. In all likelihood, both arguments will be found compelling by those antecedently disposed to thinking probability is the logic of the assessment of prospects, and found open to dispute by those who favour an alternative approach.

3.8 Conditional Degrees of Support

Frequently, one doesn't want to know only the degree of support one's current evidence provides. One wants also to know what bearing hypothetical evidence has on hypotheses one is concerned with: what are the prospects of H, *given* E, where E is not (yet) in evidence. For instance we might be interested in the support for the hypothesis that some particular die that we haven't yet tossed is fair, given the further potential evidence that it lands 6 on 100 consecutive rolls. Presumably that number will be different from the support provided by the current background evidence for its fairness, or from the support provided by the potential evidence that it lands a random mixture of 1–6 in roughly equal proportions in 100 rolls. The conditional prospect of H given E relative to some evidential probability Pr is written $\Pr(H \mid E)$.

Conditional support is related to unconditional support.[16] What is the prospect that both H and E turn out true? Intuitively,[17] it is the proportion among cases of E turning out true that are also cases where H turns out true. That is, the prospect of E turning out true, weighted by the prospect of H turning out true given E. This is captured in this rule:

Product $\Pr(H \wedge E) = \Pr(H \mid E)\Pr(E)$. (Jeffrey 2008: 12)

In line with the Product rule, the low prospect of my going to the beach *and* getting sunburnt is the prospect that I get sunburnt given I go to the beach (high), weighted by the prospect that I go to the beach (low). There is a betting-price justification of the product available, given a natural idea of a conditional bet – one that is called off when E is false (de Finetti 1937: 146; Howson and Urbach 1993: 81–84). I will not rehearse it here for reasons of

[16] Another hypothesis in this area is that conditional support is unconditional support *of a conditional*, i.e., that the degree of support of H given E is $Pr(E \to H)$ for some conditional operator. This is surprisingly hard to defend. So long as the conditional obeys Importation – that $P \to (Q \to R)$ entails $(P \wedge Q) \to R$ – then conditional probabilities will be trivial: if $\Pr(H \mid E) = \Pr(E \to H)$, then $\Pr(H \mid E) = \Pr(H)$ – E has nothing to do with it (Lewis 1976; Hájek 2011).)

[17] Thinking in terms of proportions can be misleading when thinking about probability in general (Hájek 1997, 2009), but seems safe here.

space, but none of the issues involved go significantly beyond those canvassed in §3.4.

The Product rule isn't a definition, because the right hand side mixes conditional and unconditional probability. Nor can it readily be turned into a definition, because we can only re-arrange the equation to isolate the term '$\Pr(H \mid E)$' under the assumption that $\Pr(E) > 0$. But an account of conditional probabilities as a ratio of unconditional probabilities is nevertheless often assumed, namely, the 'ratio analysis' stating that $\Pr(H \mid E) = \Pr(H \wedge E)/\Pr(E)$, provided $\Pr(E) > 0$.

Cases where the ratio analysis gives no guidance will arise commonly, since we do not insist on Regularity (§3.5); hence there will be many epistemic perspectives which give no prospect to contingent propositions. Yet in many cases of interest to us, the relevant unconditional probabilities are well-defined and non-zero. Generally, it is only important to consider the conditional support provided by E to H when E itself is genuinely in prospect. There may be some academic interest in the degree to which the hypothesis that a fair coin lands heads is supported by some independent proposition that some quantity takes a specific real-value that it has zero probability of taking (i.e., in the quantity $\Pr(\text{Heads} \mid X = x)$, where X is a continuous real-valued random variable) (Hájek 2003: 286). It is a virtue of the multiplicative relation between conditional and unconditional probability that it allows $\Pr(\text{Heads} \mid X = x) = 1/2$, even while $\Pr(X = x) = 0$ and $\Pr(\text{Heads} \wedge X = x) = 0$. But it is not especially interesting when thinking about degree of support by prospective evidence, since if the evidence has no prospect the support it may or may not provide is (relative to a given epistemic perspective) irrelevant. Of course if one finds oneself frequently confronted by evidence one regarded as having no prospect, that is probably a reason to reconsider the epistemic perspective one is deferring too; the space of possibilities it depicts is misaligned with the possibilities being actualized.[18]

[18] That said, because we won't be defining conditional prospects using the ratio analysis, it is open to us to allow conditional prospects to be defined in many cases where traditional probability theory says they are undefined. There may well be residual cases where they remain undefined; so in what follows every occurrence of generic '$\Pr(\cdot \mid \cdot)$' is to be regarded as presupposing that the conditional prospect is defined.

Suppose $\mathbf{J} = \{J_1, \ldots, J_n\}$ is a *partition* of the space of possibilities: a set of mutually exclusive and jointly exhaustive hypotheses, relative to some epistemic perspective. Given any H is logically equivalent to $(H \wedge J_1) \vee \ldots \vee (H \wedge J_n)$, both express the same proposition and have the same probability. Those disjuncts are mutually exclusive because the J_is are, so additivity entails $\Pr(H) = \Pr(H \wedge J_1) + \cdots + \Pr(H \wedge J_n)$. Apply the product rule and we get this basic result:

Total probability $\Pr(H) = \Pr(H \mid J_1) \Pr(J_1) + \cdots + \Pr(H \mid J_n) \Pr(J_n)$.

Another elementary consequence of the Product rule is that $\Pr(H \mid E) \Pr(E) = \Pr(E \mid H) \Pr(H)$. When $\Pr(E) > 0$, this can be rearranged into theorem that has a prominence belying its obvious proof:

Bayes $\Pr(H \mid E) = \dfrac{\Pr(E \mid H) \Pr(H)}{\Pr(E)}$.

Applying the theorem of total probability to '$\Pr(E)$' given a partition including H (e.g., $\{H, J_1, \ldots, J_n\}$), we can reformulate Bayes' theorem:

$$\Pr(H \mid E) = \frac{\Pr(E \mid H) \Pr(H)}{\Pr(E \mid H) \Pr(H) + \Pr(E \mid J_1) \Pr(J_1) + \cdots \Pr(E \mid J_n) \Pr(J_n)}.$$

Bayes' theorem is important because it shows how conditional probabilities $\Pr(H \mid E)$ are fixed by other quantities we often know the value of:

- The *prior* probability of the hypothesis H, $\Pr(H)$;
- The *likelihood* of the the evidence given the hypothesis, $\Pr(E \mid H)$;
- The probability of the evidence, $\Pr(E)$, which can also be expressed as the *expected likelihood* of the evidence, given hypotheses spanning the space of possibilities.

The quantity $\Pr(H \mid E)$ is often called the *posterior* probability of the hypothesis.

Here 'prior' and 'posterior' do not refer to temporal priority; $\Pr(H \mid E)$ is evaluated at the same time, and using the same evidential probability Pr, as $\Pr(H)$. There is a very common idea that *if* one is using a probability function to represent an agent's state of mind at some time, and that agent were to learn exactly E, then the agent's new state of mind should be $\Pr_E$. The rule of *Conditionalization* says that the relation between this new mental state, taking E into account, and the old mental state, is this: for any H,

$\Pr_E(H) = \Pr(H \mid E)$; i.e., that the new attitudes should be the old conditional attitudes. In light of this, the genuinely posterior state of mind $\Pr_E(H)$ is equal to a pre-existing conditional attitude, which (by extension) gets to be called the posterior. There is considerable controversy, not unrelated to our earlier discussion of reasoning in §1.3, around whether Conditionalization is required as a condition of rationality (van Fraassen 1989: 160–82; Talbott 1991; Lewis 1999; Williamson 2000: 213–21; Hedden 2015: 29–44). Again, we will set aside diachronic issues of how to update or change evidential probability functions in light of new evidence (though see §4.1). Our project, as before, is to understand evidential support given a single epistemic perspective, a single evidential probability function. Nevertheless the terms 'prior' and 'posterior' are so entrenched that it would risk intelligibility to wholly refrain from using them.

This underscores a useful clarification. A conditional degree of support need not be the degree of support you *would* assign to H if you were to find out that E. It is the degree of support H has, in light of E and your current background evidence. It could be the case that if you were to find out that E, you would revise your background beliefs. Given an epistemic perspective that represented a given die as fair, the conditional degree of support for the die landing '6' given it has landed '1' repeatedly for the past 1000 rolls is still $1/6$. But if you were to find out that regularity, you would be more than reasonable in opting to consider a different epistemic perspective, one that did not have the fairness of the die as part of its background evidence. Perhaps this makes it clear that every evaluation of prospects is really conditional in some sense, the background evidence no less important a factor in the evaluation of an 'unconditional' probability as the explicit proposition E is in the evaluation of conditional probability:

> every evaluation of probability, is conditional; not only on the mentality or psychology of the individual involved, at the time in question, but also, and especially, on the state of information in which he finds himself at that moment. (de Finetti 1974: 134)

With conditional degrees of support we have finally reached the third key role of epistemic perspectives, namely, capturing the significance of potential evidence for various hypotheses (§3.2). Given an epistemic perspective, an unconditional probability already discriminates between hypotheses in light of the the background evidence held fixed. Conditional probabilities add a new aspect to that evaluation: they give us some sense of the support various hypotheses would have given propositions not yet in evidence, where keeping the same underlying probability function is equivalent to retaining the same standards of evidential significance. So conditional and unconditional probabilities together articulate an 'epistemic standard' that might be deployed in the evaluation of hypotheses (Schoenfield 2012: 199). One simple but powerful idea that emerges from this is to pay attention to the comparative facts about the prior support H receives from the background evidence versus the posterior conditional support it receives once E is added to that background evidence, and use such comparisons to track confirmation or disconfirmation of hypotheses by E, and use comparisons between posterior probabilities for H and H' in light of E to see which is better supported by that prospective evidence. This is the key idea of *Bayesian confirmation theory* – 'Bayesian' because in comparing prior and posterior degrees of support, Bayes' theorem is often used. This is the principal focus of Section 4.

3.9 The Plurality of Epistemic Perspectives

Above we have only supposed that epistemic perspectives must satisfy the probability axioms. Very many functions do this. So if epistemic perspectives are to justify actual credences (§3.3), we shall need some way of deciding which epistemic perspective to adopt, and what exactly adopting a perspective amounts to.

The second question is easier than the first. To adopt an epistemic perspective is to take its verdicts on support as your own. If C is your credence function, representing your degrees of belief, then you adopt an epistemic perspective Pr iff for any X, $C(X) = \Pr(X)$.[19] This entails, as long as you assign credence 1 to your evidence, that you can only adopt those perspectives that have your evidence as background.

[19] The question of whether you *can* adopt an epistemic perspective is another matter – perhaps epistemic perspectives are too cognitively demanding for us. In that case, adopting an epistemic perspective is an ideal of rationality, rather than a prescriptive norm (Staffel 2020; Carr 2021).

For any given body of (consistent) evidence, there are many perspectives that assign its members probability 1. So this requirement doesn't narrow down the choice of perspective significantly. Perhaps some further requirements on perspectives can narrow them down still further? Indeed, if we could narrow it down to a single perspective, we could easily solve the problem of which perspective to adopt: one is doxastically rational in adopting Pr iff Pr is the unique epistemic perspective consistent with your total evidence.

Some suggest that it is incumbent on any rational attitude that it defer to chance (Joyce 2007-08; Ismael 2008-06; Lewis 1986a). For any quantity $\mathbf{Q}$, such that $\mathbf{Q}_i$ is a partition of hypotheses about the value of $\mathbf{Q}$, then your credence $C(H)$ will be your expectation of the value of $\mathbf{Q}$, in line with the theorem of total probability:

$$C(H) = C(H \mid \mathbf{Q}_i)C(\mathbf{Q}_i).$$

To *defer* to some quantity is for those conditional probabilities to equal the value $\mathbf{Q}$ assigns to H, when you don't have information that 'trumps' that quantity. In the case of chance that gives rise to this principle:

Chance Deference A probability P defers to chance iff for any proposition X, where '$Ch(X) = x$' states that the chance of X is x then $P(X \mid \ulcorner Ch(X) = x \urcorner) = x$.[20]

Whether a probability function defers to chance is a descriptive matter. Whether it *should* is captured by the Principal Principle, so-called because Lewis in introducing it claimed it captures 'all we know about chance' (Lewis 1986a: 86). Lewis' original formulation concerned rational credence; transferred to our context, we get something like this (remembering that epistemic perspectives are normative for credence, so if this is true, rational credences will also defer to chance):

Principal Principle: Any epistemic perspective defers to chance, on the condition that the background evidence it incorporates is *admissible* – it

[20] Here, 'Ch' denotes different probability functions at different possibilities; thus propositions of the form '$Ch(X) = x$' are contingent.

> includes no evidence that is more informative about chancy outcomes than chance itself.[21]

The Principal Principle plus the theorem of total probability will ensure that any epistemic perspective will assign a degree of support to *H* that is equal to that perspective's expectation of the chance. Adopting the Principal Principle will thus constrain the possible epistemic perspectives relative to a given body of evidence. Any probability function that takes the evidence to support *H* to a degree greater than the chance of *H*, but not because the evidence includes inadmissible evidence, will not count as an epistemic perspective. Nevertheless, we do not get a unique perspective even insisting on deference to chance, for the degree to which *H* is supported depends on the degree to which hypotheses about its chance are supported, and they are not themselves set exogenously to the perspective.[22]

One could circumvent this obstacle if it turned out that there was a unique degree to which any evidence supported hypotheses about the chance. But the only plausible scenario in which that is true is one where acceptable epistemic perspectives are already uniquely narrowed down by the background evidence, which would circumvent the appeal to the Principal Principle entirely. Again, that is the focus of Section 5.

---

21 Lewis's version concerns rational initial credence; he assumes that non-initial credence is obtained by conditionalizing on evidence, and that initial credences are regular (and hence are certain of nothing inadmissible). I make neither assumption about epistemic perspectives. Nevertheless, I do make some reference to admissibility. Others who offer versions of the PP without an admissibility clause assume tacitly that agents are certain they won't obtain inadmissible evidence (cf. Pettigrew 2020b: 6).

22 Norton (2011: 400–401) argues that some physical theories involve no chances, despite being indeterministic. He sees this as an objection to the present Bayesian framework, because he appears to take the Principal Principle as derivative from a more general maxim: that 'factual properties of the system under consideration' fix the quasi-logical structure of rational attitudes. If the system has chances, our attitudes can have a probabilistic structure; if the system is chanceless, but still indeterministic, our attitudes can only be expected to satisfy a weaker structure, perhaps comparative only rather than numerical. This objection can be resisted; epistemic rationality requires us to evaluate the prospects of hypotheses. In the absence of chances that would constrain those prospects more tightly, we must do what is nevertheless permitted, and adopt some epistemic perspective. Compare: in unrestricted sections of the Bundesautobahn system, there is no speed limit. That doesn't mean cars should somehow travel without having some determinate speed; it just means that many speeds are permissible.

Perhaps the problem of which perspective to adopt isn't to be solved by narrowing down perspectives, but by constructing a unique response to any set of perspectives. Consider again Chance Deference above; this principle has analogues for any 'expert' probability function that one is disposed to defer to (Gaifman 1988; Hall 2004: 100–101; Elga 2007: 478–80). In place of chance, one might consider an expert weather forecaster, to whom one might defer completely about a limited subject matter, or an omniscient being, to whom one ought to defer completely about everything. If 'O' is the omniscient being's credences, then substitute it for 'Ch' in Chance Deference, and one might endorse it and an appropriately modified analogue of Lewis' original Principal Principle, i.e., formulated as a constraint on credence. Again, one might not be certain what your expert says, just as you might not be certain of what the chances are; in that case, it follows by the same mathematics as above that your credence will be your expectation of the expert's opinion.

In this case it follows that you have a unique expert informed credence even when you are open to many hypotheses about the possible attitudes of the expert. The obvious move is thus open to us. If there is a single best epistemic perspective that reflects the expert judgement of how evidence supports hypotheses, then you ought to adopt as your current credence your (perhaps subjective) expectation of the extent to which your evidence supports the various hypotheses open to you. That is, you ought to have a credence in any H that equals your estimate of *the* degree to which your evidence supports H. In this way, one can arrive at a unique recommendation of which attitude to have, even if one has not come to a single conclusion on what the correct epistemic perspective is. That will give us this principle:

Deference to Evidential Probability For any H, your credence $C(H)$ should equal $\sum_i C(\ulcorner \mathrm{Pr}_i(H) = x_i \urcorner) \cdot x_i$, where '$\mathrm{Pr}_i$' ranges over those epistemic perspectives which may, for all you know, be the actual degree of support your evidence provides for H.

This rationale breaks down, however, if there are multiple epistemic perspectives consistent with the evidence. Each embodies some standards for evaluating how evidence supports hypotheses. There may be no sense in which they are rival accounts of 'the' evidential support relation. Their relation to one another may be more like the relation between the various epistemic perspectives that have different background bodies of evidence; there is no one perspective that reflects the evidence one uniquely should have, rather there are many that may reflect the various bodies of evidence one might have. Likewise, maybe there are many perspectives that may reflect the various standards of evidential support one might subscribe to. If that is the case, it is harder to justify adopting as your credence the weighted mean of these various epistemic perspectives. This is because we can only rationalise expert deference if there is an expert, if there is a fact of the matter about getting degree of support right that transcends each individual evidential probability function.

Without a unique degree of support, there may be no way to recommend a unique attitude in view of the plurality of epistemic perspectives. That will be bad news for a restrictive epistemology, which aims to tell us what we must do. But if epistemology is *permissive* – in the business of telling us what we *may* do – then a natural idea suggests itself: adopt any epistemic perspective consistent with your evidence. As van Fraassen puts it: 'rationality is only bridled irrationality … what it is rational to believe includes anything that one is not rationally compelled to disbelieve' (van Fraassen 1989: 171–72). This suggests:

Permissive Rational Credence There are many epistemic perspectives Pr_1, Pr_2, etc. An agent's credence function at a time, C, is rational just in case there exists a probability function Pr_i such that for any H, $C(H) = \mathrm{Pr}_i(H)$, where Pr_i agrees with C on the total evidence (cf. Climenhaga 2023: 6 for a related idea).[23]

Because each epistemic perspective embodies a conception of evidential relevance,

Ultimately, then, the issue of rational credence boils down to the issue of uniqueness of epistemic perspectives. That, as mentioned above, is the focus of Section 5. So far we have found no reason to accept that there a single perspective compatible with each body of total evidence; even if there are constraints on perspectives beyond the probability axioms – such as the Principal Principle – they appear to leave open many legitimate attitudes to evidential

[23] There are many varieties of permissivism, depending on just what is permitted (Meacham 2014: 1186–89). Our sympathetic discussion of the Principal Principle suggests a moderate rather than extreme permissivism.

support, as represented in the existence of many evidential probability functions that agree on the total evidence. Before returning to that issue, however, I return in Section 4 to our central question: applying the epistemic perspective framework to our motivating puzzle of inductive support.

Section 4

Bayesian Confirmation Theory

In Section 2 we reviewed various formal and syntactic approaches to the logic of evidential support, and found them wanting. In Section 3 we suggested that a better approach to evidential support was through the idea of the prospects of hypotheses in light of a body of background evidence and assumptions about the bearing of evidence on hypotheses, jointly encapsulated in the notion of an epistemic perspective, normative for rational credence. We argued that epistemic perspectives ultimately were to be identified with evidential probability functions, and that those functions enable us to define a notion of overall evidential support of a proposition.

In this section, we look at the incremental support provided by a piece of evidence, the relation of confirmation, defined and applied to some representative cases in §4.1. I turn immediately to a famous problem for the Bayesian approach to confirmation, the problem of old evidence (§4.2), and argue that it is not an insuperable objection to our approach. §4.3 introduces a related probabilistic approach to confirmation, Likelihoodism, and illustrates a key difference between it and the Bayesian approach on the issue of base rates and their evidential role. Further illustrations of successes of the Bayesian approach in capturing scientific heuristics are given in §4.4, and we return to the paradoxes of confirmation in §4.5. The fullest Bayesian treatment of the paradoxes requires a notion of degree of confirmation, and there are a number of inequivalent attempts to measures that, with varying degrees of plausibility; some are introduced and compared in §4.6. I conclude by returning to the problem of induction (§4.7).

The Bayesian approach, despite its flexibility and adaptability, has not avoided criticism. I will touch on some significant objections as the view is developed, and have hoped to forestall some by design. Some prominent criticisms focus on aspects of some Bayesian views that are not representative of the variant I have defended. Consider challenges based on the supposed computational intractability of Bayesian statistics, or the failure of scientists to assign or report probabilities in practice, or the descriptive inaccuracy of the Bayesian picture (Kelly and Glymour 2004: 95–96). These may be telling against personalist Bayesianisms, which take the probabilities involved in support and confirmation as unprocessed credences. But it is hard to see how they are to apply to the degree of support picture, which posits evidential probability as a regulative ideal, perhaps not directly implemented in practice. Other objections are more general; Norton (2011: 400–415) is keen to emphasise ways in which the Bayesian framework appears to require a richer structure on our attitudes than is sometimes justified by the evidence (recall §3.9). That could be a problem if the Bayesian is committed to uniqueness (for then there must be a uniquely rational attitude ungrounded by evidence sufficient to constrain the attitudes uniquely). It appears to be less of a problem for the permissive Bayesian, who allows that there are many legitimate perspectives that can be taken without being uniquely determined – what matter, then, if we are permitted to believe beyond what the evidence demands in these cases too? So this family of objections turns out to connect with what I regard as the principal objection to Bayesianism, the 'problem of the priors'. This objection is the focus of Section 5.

4.1 Incremental Confirmation Defined

Each evidential probability/epistemic perspective carries with it some conception of background evidence. Concerning a proposition E not yet in evidence, we can use the conditional evidential probability to indicate the degree of support that H has in light of the background evidence supplemented with E (§3.8). We introduce the notion of qualitative *incremental* evidential support, or *confirmation*, as a three-place relation between some piece of *potential* evidence E, some hypothesis H, and a background probability model which is assumed to represent an epistemic perspective and hence to capture both background knowledge and a specific conception of evidential relevance (cf. Fitelson 2005: 391).

Confirmation E *confirms* H relative to Pr iff $\Pr(H \mid E) > \Pr(H)$.
E *disconfirms* H relative to Pr iff $\Pr(H \mid E) < \Pr(H)$.
E is *independent*[1] of H relative to Pr iff $\Pr(H \mid E) = \Pr(H)$.[2]

This definition is at the heart of *Bayesian confirmation theory* (Earman 1992; Howson and Urbach 1993; Strevens 2006). The term 'Bayesian' is apt because, given this definition of confirmation and Bayes' theorem, E confirms H iff $\frac{\Pr(E|H)}{\Pr(E)} > 1$. Given that $\Pr(E)$, the degree to which the background evidence supports the evidence, is constant for any hypotheses we may consider, this result tells us that the confirmation of H by E is driven by the likelihood $\Pr(E \mid H)$, and this is a quantity we are often in a good position to determine, because it is often fixed by the content of H itself. Suppose H is a chancy theory, for example; then the Principal Principle (§3.9) tells us that the likelihood of the evidence on H is its chance of coming about were H the correct theory of E's chance (Strevens 2004: 372–74). Despite the explicit relativisation of evidential support to a probability model, whether a hypothesis H is confirmed turns out to depend on the likelihoods alone, which are in many scientific cases of interest fixed by H itself, and thus common to many epistemic perspectives. What still varies from perspective to perspective is the absolute degree of support for a given hypothesis, whether that hypotheses for example reaches some minimal level of credibility in light of the evidence.

Let's consider a toy example.

Biased Die Suppose we have background information that leaves open the hypotheses that a die is biased towards '6' (the chance of '6' being 1/3), or that it is fair, and regards these hypotheses as equally supported and exhaustive, so Pr(fair) = Pr(6-biased). No outcomes are yet in evidence.

1 There are some complications here (Fitelson and Hájek 2014); note too that this notion of probabilistic independence neither implies nor is implied by other sorts of independence such as causal isolation.

2 This account of confirmation goes back at least to Carnap (1962: xvi), where Carnap distinguishes 'increase in firmness', or confirmation in our sense, from 'firmness', which is the degree to which some evidence supports a hypothesis. In light of our earlier discussion in §3.8, the degree of support is just the conditional probability $\Pr(H \mid E)$, which involve no relation to any other degree of support.

Against this background, consider the proposition that the die lands '6' 5 times consecutively. That proposition confirms the hypothesis that the die is biased, disconfirms the hypothesis that the die is fair.

Can Bayesian confirmation theory capture this basic observation about confirmation? Eventually, we might use a theory of confirmation to decide complex cases about which our judgements are unclear, but as usual with a philosophical explication, we need reassurance that it gets the basics right, typically by checking that it reproduces obvious facts about the target relation. In this case, we begin with an appeal to the theorem of total probability to calculate the prior of E:

$$\begin{aligned}\Pr(E) &= \Pr(E \mid \text{fair}) \Pr(\text{fair}) + \Pr(E \mid \text{6-biased}) \Pr(\text{6-biased}) \\ &= \frac{1}{6^5} \cdot \frac{1}{2} + \frac{1}{3^5} \cdot \frac{1}{2} \\ &= \frac{1}{15,552} + \frac{1}{486} \approx 0.002122.\end{aligned}$$

Plug this into Bayes' theorem (§3.8), and we obtain

$$\Pr(\text{fair} \mid E) \approx 0.0606\,\mathrm{Cr}(\text{fair}) < \mathrm{Cr}(\text{fair});$$
$$\Pr(\text{6-biased} \mid E) \approx 1.9394\,\mathrm{Cr}(\text{6-biased}) > \mathrm{Cr}(\text{6-biased}).$$

So here this proposition confirms the hypothesis of bias, and disconfirms the hypothesis of fairness.

The same sort of reasoning can be applied in less idealized examples. Consider Howson and Urbach's account of Babbage's investigation into the origins of tables of logarithms:

Logarithms 'Babbage … was interested in whether they derived from the same source or had been worked out independently. Babbage (1827) found the same six errors in all but two and drew the "irresistible" conclusion that, apart from these two, all the tables originated in a common source.' (Howson and Urbach (1993), p. 124)

Howson and Urbach offer this Bayesian reconstruction (following Jevons (1874), pp. 278–9):

> The theory Copy, which says of some pair of logarithmic tables that they shared a common origin, is moderately likely in view of the immense amount of labour needed to compile such tables *ab initio*, and for a number of other reasons. The alternative, independence theory might take a variety of forms, each attributing different probabilities to the occurrence of errors in various positions in the table. The only one of these which seems at all likely would assign each place an equal probability of exhibiting an error and would, moreover, regard those errors as more-or-less independent. Call this theory Ind and let E^i be the evidence of i common errors in the tables. The posterior probability of Copy is inversely proportional to $\Pr(E^i)$, which, under the assumption of only two rival hypotheses, can be expressed as $\Pr(E^i) = \Pr(E^i \mid \text{Copy}) \Pr(\text{Copy}) + \Pr(E^i \mid \text{Ind}) \Pr(\text{Ind})$. … Since Copy entails E^i, $\Pr(E^i) = \Pr(\text{Copy}) + \Pr(E^i \mid \text{Ind}) \Pr(\text{Ind})$. The quantity $\Pr(E^i \mid \text{Ind})$ clearly decreases with increasing i. Hence $\Pr(E^i)$ diminishes and approaches Pr(Copy), as i increases; and so E^i becomes increasingly powerful evidence for Copy, a result which agrees with scientific intuition. (Howson and Urbach 1993: 124–25, notation adjusted)

This reasoning contains two interesting observations. First, the observation that when a hypothesis entails some evidence, $\Pr(E \mid H) = 1$ (by the probability calculus). Then $\Pr(H \mid E) = \Pr(H)\frac{1}{\Pr(E)}$; unless the probability of the evidence is 1, therefore, evidence entailed by a theory supports it. Second, in this case it is argued that $\Pr(E^i)$ approaches $\Pr(H)$ as the number of common errors goes up (because of the decreasing chance of such a coincidence). So not only does E^i confirm H; the degree of support E^i provides for H approaches 1. We see here not only an invocation of comparative relations of confirmation, but 'absolute' degree of support in light of the evidence. An epistemic perspective with this kind of conception of the available hypotheses, and these ideas about likelihoods informed by background assumptions about chances, supports the proposition that the common errors arise from plagiarism, not luck; hence any rational attitude adopting such a perspective must inevitably reach the same conclusion.

It is tempting to many Bayesians to import their views on diachronic rationality into confirmation theory. From this perspective, 'confirmation' is a thing that happens to a theory when new evidence arrives; the theory is confirmed or disconfirmed as its probability shifts around over time. Most Bayesians accept updating upon receipt of new evidence E goes by conditionalization, adopting one's old conditional credences given E as one's new 'unconditional' credences. In this case, $\Pr(H \mid E)$ represents the new posterior probability of H; H is confirmed by E if it has come to be more probable once news of E is in. So one sometimes sees Bayesians present confirmation as essentially diachronic:

> an experience provides evidence that confirms a hypothesis, for that scientist, if … this evidence 'boosts' the scientist's credence in the hypothesis. (Eagle 2011: 210)

> At the core of modern Bayesianism is a rule for changing the subjective probabilities assigned to hypotheses in the light of new evidence. … where $C(\cdot)$ is your subjective probability distribution before observing E and $C^+(\cdot)$ is your subjective probability distribution after observing E, …
>
> $$C^+(H) = \frac{C(E \mid H)}{C(E)} C(H).$$
>
> More or less anyone who counts themselves a proponent of BCT thinks that this rule is the rule that governs the way that scientists' opinions should change in the light of new evidence. (Strevens 2004: 369; see also Strevens 2006: §5.1)

But this really runs together two completely separate issues: whether conditionalization is the right updating rule, and under what synchronic circumstances does one proposition support another. (Recall here §§1.3, 3.8.) It is perfectly possible to endorse the Bayesian account of evidential support while rejecting conditionalization. Better, then, to interpret the confirmation inequality $\Pr(H \mid E) > \Pr(H)$ as telling us the current bearing of E on H from the single perspective of Pr, a stance that may or may not be reflected in a relation between successive perspectives, pre- and post-acquisition of E.

4.2 Old Evidence

This bears on a problem which has been seen by many as a serious challenge to the Bayesian account of confirmation, Glymour's problem of *old evidence*

> Scientists commonly argue for their theories from evidence known long before the theories were introduced. Copernicus argued or his theory using observations made over the course of millennia, not on the basis of any startling new predictions derived from the theory, and presumably it was on the basis of such arguments that he won the adherence of his early disciples. Newton argued for universal gravitation using Kepler's second and third laws, established before the *Principia* was published. The argument that Einstein gave in 1915 for his gravitational field equations was that they explained the anomalous advance of the perihelion of Mercury, established more than half a century earlier. ... Old evidence can in fact confirm new theory, but according to Bayesian kinematics it cannot. For let us suppose that evidence E is known before theory T is introduced at time t. Because E is known at t, $\Pr_t(E) = 1$ [so] the likelihood of E given T, $\Pr_t(E \mid T)$, is also 1. We then have
>
> $$\Pr_t(T \mid E) = \frac{\Pr_t(T) \times \Pr_t(E \mid T)}{\Pr_t(E)} = \Pr_t(T).$$
>
> The conditional probability of T on E is therefore the same as the prior probability of T: E cannot constitute evidence for T in virtue of the positive relevance condition nor in virtue of the likelihood of E on T. None of the Bayesian mechanisms apply, and if we are strictly limited to them, we have the absurdity that old evidence cannot confirm new theory. (Glymour 1981: 85–86)

As many have noted, the problem is not so much the antiquity of the evidence, as the fact that evidence and hypothesis seem to come in the wrong temporal order. If the age of the evidence were the only problem, we could solve the problem by 'rolling back' to an earlier state of knowledge in which the crucial evidence isn't included – this seems to be what Howson and Urbach have in mind when they say '$\Pr(H)$ measures your belief in a hypothesis *when* you do not know the evidence' (1993: 117, my emphasis). But Climenhaga (2023: §3) sets up a simple example in which a piece of evidence is acquired, then some crucial information about the probability distribution over hypotheses is acquired, and then a posterior probability over hypotheses is calculated. Roll back the credence to before the acquisition of the evidence, and one also loses the distributional information. There was never, in his case, a state of belief that represented the background against which this evidence is confirmatory in the way it appears to be.

Introduced like that, old evidence is not a problem for the Bayesian view I have presented, which involves no 'kinematic'/diachronic element. Though I follow the Bayesian literature in talking of confirmation of hypotheses by evidence, and use the suggestive variable 'E', I explicitly reject the idea that confirmation occurs when evidence is newly acquired.[3] There is no sense in which the evidence considered with respect to confirmation has to be collected at all. Recall our discussion in §3.2; an epistemic perspective is associated with some body of total evidence, some body of propositions (it turned out) that are assigned probability 1 by the perspective. Nothing temporal is involved in this characterisation. Utilize an epistemic perspective including E as evidence, and it won't confirm anything; utilize an epistemic perspective not including E, and it may well have confirmatory power. This is independent of when E is gathered, or even if it is gathered. The incremental confirmation relation informs us of the evidential bearing of one proposition on another, relative to an epistemic perspective; neither needs to be 'evidence' in a folk or philosophical sense for this evidential bearing to obtain. Rather, E is some claim that might bear on H – perhaps H predicts it, or some rival of H predicts it – and we wish to evaluate its significance, relative to some perspective embodying some appropriate principles of evidential bearing.

To do this, of course, one must make use of an epistemic perspective according to which there is some bearing of E on H. Glymour's argument certainly emphasises that a perspective that assigns probability 1 to E is not appropriate for this purpose. Nor for that matter is one that assigns probability

3 Here I claim illustrious precedent (Hempel 1945a: §6; Carnap 1962: 468).

1 to H, which would then be incapable of being confirmed. Nothing in the framework I've present requires us to make such inappropriate choices; but nothing tells us which choices to make, either. So the *synchronic* problem of old evidence is to give defensible guidance about which epistemic perspective should we consider when we evaluate confirmation of hypotheses by claims which are already in evidence for us.

One obvious candidate is clearly excluded, because it will simply reinscribe the problem of old evidence. This is the proposal that we ought to evaluate claims confirmation relative to an epistemic perspective we've adopted. Our theory of adoption in §3.9 guaranteed that anything in evidence for us is assigned probability 1 by any adoptable perspective. Hence the problem of old evidence shows that even at a fixed point in time there is no single body of background evidence: the background evidence relevant to the adoption of an epistemic perspective is the evidence possessed by the adopting agent, which may be different than the body of evidence against which that very same agent assessed claims of confirmation. A theory of confirmation is quite distinct from a theory of individual belief, as Glymour (1981: 74) pointed out; this unworkable proposal would collapse them.[4]

Some suggest simply removing E from the background evidence, and evaluating all confirmation claims relative to an epistemic perspective that treats that background (mutilated, from our perspective) as its total evidence:

> One answer-and I think the correct one-to Glymour's nasty problem … is to deny that when assessing support according to the difference between $\Pr(H \mid E)$ and $\Pr(H)$, the probabilities should be relativized to K; rather they should *always* be relativized to $K \setminus \{E\}$ … And why? The answer is straightforward. When you ask yourself how much support E gives H, you are plausibly asking how much a knowledge of E *would* increase the credibility of H, which is the same thing as asking how much H boosts the credibility of H relative to what *else* you currently know. The 'what else' is just $K \setminus \{E\}$. (Howson 1991: 548)

4 Glymour assumes that the Bayesian can only be giving a theory of individual belief, but our theory is obviously not that, given that it needs bridge principles of some degree of controversy to link actual credence to epistemic perspectives.

The proposal is an example of a more general class of *counterfactual* theories, those that evaluate the confirmatoriness of E with respect to how much E 'would increase' the credibility of H against a counterfactual background, what we would have known had we not known E. Howson's approach is, in effect, that we would have known everything but E. Howson's suggestion seems to give incorrect predictions about confirmation, of two sorts.

Sometimes a generalisation is confirmed by its instances (§4.5), and, in the Baconian fashion (Bacon 1620), may only be proposed after diligent apian collection of facts. We know that there is a point at which confirmatory returns to repeated experiments must diminish to zero (Howson and Urbach 1993: 120). In such a case, subtracting any particular instance from background knowledge leaves many similar instances, and no confirmation by that instance. The same is true, clearly, for every instance – so no instance confirms. Howson seems blazé about this (1991: 550). Yet intuitively each instance may be highly confirmatory were *none* of the similar instances present, and our instincts in particular cases go with this latter observation. In this case the counterfactual about what the Baconian would have believed had they not believed E seems to give us $K \setminus \{E\}$, as Howson claims, but that seems to be the wrong body of background evidence to use in evaluating confirmation.

In this case, the counterfactual yields the 'wrong' body of background evidence to assess confirmation. In other cases, the counterfactual yields the right background, but the wrong epistemic perspective: it not only subtracts E from the background knowledge, it also shifts us to an epistemic perspective in which evaluations of the bearing of E on H are different. Maher gives this example:

> Mr. Schreiber is the author of novels that are popular (P) though it is important to him that he is making important contributions to literature (I). Schreiber basks in his success, taking his popularity to be evidence of the importance of his work; that is, he takes P to confirm I. … many aspiring serious novelists whose work is unpopular tend to rationalize their failure by supposing that the public taste is so depraved that nothing of true value can be popular. … if Schreiber did not know of his own work's popularity, he too would share this opinion … [That is,] were he not to know P, he

would have a probability function Pr such that $\Pr(I \mid P) \leq \Pr(I)$. (Maher 1996: 156)

What we want is something like this: a surgical modification of our current adopted epistemic perspective that preserves our dispositions to evaluate the bearing of E on H,[5] while removing E and evidence substantially similar to E to predict the right judgments about confirmation. That suggests the following proposal (Eells and Fitelson 2000: 667–69; Jeffrey 2008: 44–47; Meacham 2016: 461–62):

Ur-Probability In an assessment of confirmation, the probability function Pr must be such that

1. There is some *ur*-probability $\Pr_0$ that does not assign unconditional probability 1 to any proposition that confirms or disconfirms H, such that, where V is one's current total evidence, $\Pr_0(\cdot \mid V)$ is an adoptable epistemic perspective;[6]
2. There exists some 'contextually determined background evidence' B (Meacham 2016: 462), such that $B \subseteq V$ and $\Pr(\cdot) = \Pr_0(\cdot \mid B)$.

This proposal, unlike Howson's, doesn't require the determinacy of any counterfactual claim about what attitudes we might have had supposing we had different evidence. The relativity to background evidence is made explicit in a way that permits us to remove more than E if needed; but we preserve judgements of evidential bearing, by requiring that the *ur*-probability be one that could end up with an adoptable evidential probability.

The role of background is vital, because of another flaw in Howson's approach we haven't yet noted: namely, very often, old news is no news. The example of perihelion of Mercury is rather unusual; many pieces of evidence have no confirmatory value at all, being so thoroughly absorbed into our perspectives on the world that no context renders them as foreground. It is perfectly reasonable to think that the fact that something exists, for example, is

5 Though of course those dispositions are currently not well represented by our current conditional credences, which are trivialised because of the old evidence on which we condition.

6 This also solves the problem of old theories; theories of which we are almost certain cannot have their degree of support increased much even by strong evidence; but we can still analyse this notion of 'strong evidence' relative to a perspective in which the hypothesis is not highly probable.

part of my total evidence, and forms part of the base of support for various hypotheses I entertain. It would be very strange to think of this fact as confirming any hypotheses, since no live hypothesis is incompatible with it.[7] So it is important to note that our proposal above does not require that $E \notin B$, though obviously in many cases it will be included.

This proposal tells us that whether, and to what extent, E confirms H is relative to background assumptions in two ways. First, it is relative to some assumption about evidential bearing, encoded in the prior conditional probabilities. Second, it is relative to some selection of background evidence. Neither of these relativities collapse into one another. If the thesis of Uniqueness is true (§3.9, Section **??**), then there is only one legitimate perspective on evidential bearing, and yet there remain many possible selections of background evidence. On the other hand, one could accept permissivism about what bears on what, and think that epistemology proper must always consider the current total evidence in evaluating the justification of belief. The problem of old evidence brings the second sort of relativity to background into clear view: even at a given point in time, there is no single body of evidence that it is pertinent to confirmation.

Can we be more specific about how context selects background evidence? General Gricean principles governing conversation are more helpful here than special-purpose considerations about theory testing (Grice 1989). Rather than multiply theoretical posits beyond necessity, I will simply identify the background evidence with a conversational context: for example, that of a discussion between scientists about the merits of a theory, in person, or in the pages of the journals. (It could even be a private conversation the individual has with themselves.) A context, for Stalnaker, comprises 'the body of information that is presumed, at that point, to be common to the participants in the discourse' (Stalnaker 1998: 98). There are a couple of ways to think about old evidence given this. One is to note the diversity of the

7 Proponents of the so-called 'Fine-Tuning Argument' for Theism dispute this, favouring the use of epistemic perspectives that are not certain that the universe permits life to exist (Barnes 2020: 1225), though no such perspective is adoptable. Such perspectives certainly exist, but their practical relevance to theory confirmation is more dubious than proponents of FTA appear to acknowledge.

scientific community; perhaps the scholarly community isn't all of a uniform opinion about the perihelion of Mercury, for example.[8] In that case this old evidence, though known to some, cannot be common ground, and hence will not be included in the background evidence.

More interesting is a second approach. Suppose you and I are talking about a brand new hire at our company, Jack, with whom I've just had an unpleasant interaction. I ask 'why is Jack so irritable?', and you say, 'He's just stopped smoking'. I don't know Jack at all; it is not common ground to us that Jack was previously a smoker. The sentence however *presupposes* that Jack used to smoke – saying it is only legitimate on that assumption, so perhaps it ought to misfire if that assumption isn't common ground. That is does not is, Lewis suggests, due to *accommodation*:

> If at time t something is said that requires presupposition P to be acceptable, and if P is not presupposed just before t, then – *ceteris paribus* and within certain limits – presupposition P comes into existence at t. (Lewis 1979: 340)

That we accommodate the conversational moves of other speakers, I contend, also makes sense of conversational contributions that force the *retraction* of items in the common ground. Because the common ground is a set of assumed propositions, any conversational contribution that expands the possibilities under consideration can remove propositions from common ground. This, Lewis thinks, is how the sceptical argument works. What we know is what holds in all possibilities consistent with our evidence. If some new possibility is raised to salience, then our evidence is revised – weakened – so as not to exclude the new possibility. This is Lewis' 'Rule of Attention' (Lewis 1996: 559): when you attend to some new possibility, you do not ignore it, and hence do not know it not to obtain, and hence it can no longer be common ground. This also looks like a kind of accommodation: 'possibly, P' presupposes that $\neg P$ is not in the common ground; if a speaker says it, we accommodate their utterance by ensuring that $\neg P$ *isn't* in the common ground. (The

8 Or perhaps (more likely) they aren't all uniform in their grasp on how that piece of evidence is predicted by the theory (Garber 1983) – the derivation will be an auxiliary piece of old evidence, though this may run up against the questions concerning omniscience about necessities that we've tried to set aside in our treatment of epistemic perspectives in §3.2.

same is true if a speaker says P which we had been assuming was common ground. Stalnaker (1973: 454) says 'it is in general required that the proposition which is expressed by the use of a sentence in a context *not* be presupposed in that context'; if this general requirement is imposed, to say P demands – and automatically receives – accommodation if your conversational partners were assuming $\neg P$.)

These sorts of examples, I contend, might play a role in confirmation by old evidence. Ask questions about confirmation, such as 'is the theory of relativity supported by the evidence?', 'does the geological record confirm Uniformitarianism?', etc., and we trigger an accommodatory shift to a context in which the common ground includes neither evidence nor hypothesis. If it did, those questions would generally be trivially answerable by the questioner already. A broadly Gricean account then says: the questioner must not be presupposing what they are asking about, on the assumption that they are being cooperative. And then the context shifts to accommodate the speaker's not presupposing E by ensuring that $\neg E$ becomes (again) a live possibility, against which conversational background we can then evaluate $\Pr(H \mid E)$ so that it is a non-trivial question whether it exceeds $\Pr(H)$.

There is doubtless more to say about how background context selects a relevant body of evidence against which confirmation is evaluated. But saying it is not specifically the job of a logic of confirmation. And what has been said suffices, I think, to fend off the challenge of old evidence.

4.3 Likelihoods and Contrastive Confirmation

Some students of the scientific method analyse the cases from §4.1 slightly differently. They are sympathetic to the idea that probability is central to evidential support. But they are suspicious of the apparent arbitrariness that goes into selecting a particular epistemic perspective to assign prior probabilities to evidence and hypothesis (Royall 1997: xiii),[9] and unconvinced by any of the proposals we will discuss later (Section 5) to allow the background evidence to constrain that selection uniquely. However, these philosophers are impressed with the apparent objectivity of the likelihoods that play such an important

9 The evidential probability framework ameliorates to a certain extent this charge of 'subjectivity': there is nothing subjective about confirmation relative to a probability model.

role in confirmation. They want to rest their whole account of confirmation on likelihoods, and set aside prior probabilities of evidence and hypothesis as much as they can. Let us take a closer look at these *Likelihoodists* (Sober 1994; Milne 1996; Royall 1997).

A proposition E supports H when $\frac{\Pr(H|E)}{\Pr(H)} > 1$. This might lead us to define a contrastive notion of when some single piece of evidence *favours* one hypothesis over another:

Contrastive Favouring (Ratio) E favours H over H' iff

$$\frac{\Pr(H \mid E)}{\Pr(H)} > \frac{\Pr(H' \mid E)}{\Pr(H')}.$$

This definition gets the suffix 'Ratio' because it invokes a particular measure of *how much* evidence supports a hypothesis that grounds the contrastive claim; in this case, the 'Ratio measure' (Fitelson 2007: 478). We will return to measures of confirmation below in §4.6.

This contrastive notion can hold even when E doesn't confirm H; the ratio of posterior given E to prior for H doesn't exceed 1, but exceeds that for H'. For instance, tossing a die ten times and landing all '6's disconfirms the hypothesis that the die is fair, and disconfirms the hypothesis that the die is biased towards '1'. But in disconfirming both hypotheses, nevertheless it favours the 'fair' hypothesis over the biased-to-'1's hypothesis, because it doesn't disconfirm the former as much as it does the latter.

By Bayes' theorem,

$$\frac{\Pr(H \mid E)}{\Pr(H)} = \frac{\Pr(E \mid H)}{\Pr(E)}.$$

So it follows from Contrastive Favouring (Ratio) that E favours H over H' iff $\Pr(E \mid H) > \Pr(E \mid H')$. This contrastive notion reduces to a pure comparison of likelihoods; this is therefore a Bayesian derivation of the

Law of Likelihood Evidence E favours hypothesis H over hypothesis H' iff $\Pr(E \mid H) > \Pr(E \mid H')$ (Hacking 1965: 106–9; Royall 1997: 3; Gandenberger 2016: 4–5).

Likelihoodists make this Law the central plank of their theory of confirmation. They want to set aside absolute confirmation, and indeed absolute degree of support, and corresponding questions about belief and action (Royall 1997: 4). As Milne puts it,

> Confirmation concerns the support given to hypotheses by evidence. Let us suppose that two theories both entail evidence statement E and E is found to be true. The fact that E is the case is then powerless to discriminate between the two hypotheses: both entail it so there is nothing *in the nature of the evidence itself* that yields grounds on which to differentiate between them. There may be any number of *other* reasons, such as simplicity, explanatory power, and parsimony of ontological commitments, for preferring one to the other but E's truth in and of itself provides none. (Milne 1996: 22)

Our notion of the conditional degree of support E provides to H, $\Pr(H \mid E)$, runs directly against this Likelihoodist stricture. For differences the unconditional degree to which background evidence supports H as against H' can make a difference between the support E gives to H and H', even when both entail E.

The Likelihoodist approach is only plausible if there is a firm division between the 'public' and 'better known' likelihoods (Hawthorne 2005: 278) and the 'private' unconditional degrees of support used by Bayesians. If the priors are objectively given, the Likelihoodist ought to use them (as Likelihoodists accept: Sober (2008), p. 32). If the likelihoods are not objectively given, then Likelihoodism doesn't apply (or runs the risk of hypocrisy). This second horn is a real threat. In some cases the hypothesis of interest is an 'umbrella hypothesis', a statistically complex mixture of hypotheses each of which supports a precise probability distribution, but where there is no objective weighting to determine the contribution each distribution makes to the overall likelihood function (Gandenberger 2016: 5–6). For example, suppose our die example involved the hypotheses 'fair' and 'biased', where the latter is the umbrella hypothesis including many types of bias. The likelihood the hypothesis of bias assigns to a given sequence of rolls will heavily depend on the relative contribution to that likelihood of the different likelihoods fixed by different biases. A sequence of all '1's may be likely or unlikely on the hypothesis of bias, subject to the prior probability of the hypotheses 'biased towards 1s'. The principled Likelihoodist should concede: talk of evidential favouring should be avoided in such cases, regardless of how common such cases are in scientific practice.

Focusing on statistically simple hypotheses, let's consider how this might work in practice.

Mammogram Mary, a woman aged 47, attends her annual mammogram. Mary has no other symptoms, but while she is waiting for the initial results, she worries: what if there is something abnormal on the scan? Won't that be evidence of cancer?

Mary has some reason to worry, but perhaps not as much as we might antecedently think. Suppose we consider an epistemic perspective that incorporates information about the error rates for mammograms for women in Mary's situation.[10] Table table 4.1 shows some relevant aggregate statistics about mammogram results.

Table 4.1: Aggegate test results for 1,682,504 mammograms 2007–2013 (Lehman, Arao, *et al.* 2017: 53, table 2).

	Normal mammogram	Abnormal mammogram	Total
No cancer	1,486,553 (TN)	186,139 (FP)	1,672,692 ($\neg C$)
Cancer	1,283 (FN)	8,529 (TP)	9,812 (C)
Total	1,487,836 ($\neg A$)	194,668 (A)	1,682,504

The true positive rate, otherwise known as the *sensitivity* of a test, is the likelihood of an abnormal mammogram given cancer. Given these statistics,

$$\Pr(A \mid C) = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}} = \frac{8529}{9812} = 0.869.$$

10 The actual data is complex: there are many different screening regimes, with different error rates, and different ways of categorising the relevant population. The following probabilities are broadly indicative.

The true negative rate, or *specificity* is the likelihood of a normal mammogram given that the test subject doesn't have cancer:

$$\Pr(\neg A \mid \neg C) = \frac{\text{TN}}{\text{TN} + \text{FP}} = \frac{1,486,553}{1,672,692} = 0.889.$$

Accordingly, the likelihood of an abnormal mammogram given no cancer is $\Pr(A \mid \neg C) = 1 - \Pr(\neg A \mid \neg C) = 0.111$ (cf. Hendrick and Helvie 2011: W113). These likelihoods, given the Law of Likelihood, tell us (unsurprisingly) that an abnormal mammogram is evidence favouring cancer over no cancer (at the time of screening). This is also in accordance with the Bayesian framework.[11]

Having noted that the positive test is evidence for cancer, the Likelihoodist regards the job of the biostatistician as complete:

> our objective as statisticians is to understand how the data should be presented and interpreted as evidence about the risks … The published paper presents the data, along with analyses that make clear its evidential meaning. The readers will then use the evidence to adjust their beliefs and to help them in making decisions. (Royall 1997: 4)

But it would be a poor idea in Mary's case to 'adjust her beliefs': she ought not adjust her beliefs about cancer – not even comparative beliefs – on receipt of evidence about favouring. On this data, the cancer rate among all screenings is $\Pr(C) = 9812/1,682,504 \approx 0.006$. This is the *base rate*. Bayes' theorem tells us that the predictive value of a positive test result is not the sensitivity of the test, but rather the probability of cancer given an abnormal result:

$$\Pr(C \mid A) = \frac{\Pr(A \mid C)\Pr(C)}{\Pr(A \mid C)\Pr(C) + \Pr(\neg A \mid C)\Pr(\neg C)} = \frac{0.869 \cdot 0.006}{0.869 \cdot 0.006 + 0.111 \cdot 0.994} \approx 0.045.$$

So while we see incremental confirmation of cancer from the abnormal result, we see that the degree to which the total evidence supports a cancer diagnosis is very low.

[11] If the partition of hypotheses is $\{H, \neg H\}$, and E confirms H, then E disconfirms $\neg H$; so in this sort of case, E favours H over $\neg H$ iff E confirms H and disconfirms $\neg H$. We've already seen this isn't generally true for pairwise comparisons of hypotheses that are not jointly exhaustive.

This may be a problem. The Likelihoodist wants to separate questions of what is evidence for what, from questions about individuals might respond to the evidence. Having separated them, however, it turns out that questions about which hypotheses are to be accepted or relied on fall on the individual side, and aren't the kind of thing science is in the business of evaluating.[12] This seems, to many, tantamount to giving up on science – a science that isn't about recommending hypotheses for acceptance is not worth doing (Gandenberger 2016: §3):[13]

> if confirmation doesn't have a close connection with learning, that only undermines its importance – for the main aim of scientific inference isn't to see what confirms what for its own sake, but to discover what we should believe. (Eagle 2011: 215; see also Brössel and Huber 2015: 740)

This sort of stance is clearly tempting to some Likelihoodists too; for example, Barnes (2020: 1225) offers a Likelihoodist version of the Fine-Tuning Argument for theism that concludes 'the existence of a life-permitting universe strongly favours theism over naturalism'; but of course, on Likelihoodist precepts, this cannot be an argument for the existence of God, but only a claim about the comparative probabilities of the data. The posterior probability of God's existence could still be arbitrarily low and this Likelihoodist conclusion be true; but in presenting this as an argument for theism, Barnes has been tempted into an illegitimate (given Likelihoodism) claim about posterior probabilities.

Despite the gung-ho attitude of my earlier self, the lack of connection with belief and action *cannot* be decisive, given our approach to confirmation. For if confirmation is understood as a three-place relation, requiring the specification of a probability model, then any connection with belief or action must

12 The Likelihoodist may not include the base rate in their account of evidential support, but it is wholly consistent with Likelihoodism to make consideration of the base rate mandatory in the management of belief. This contrasts with the situation in frequentist statistical hypothesis testing (significance testing), where it has been argued that neglect of the base rate is encouraged by the frequentist denial that probabilities can be meaningfully assigned to hypotheses, with problematic consequences (Howson 2000: 54).

13 This is also the basis for a certain kind of response to the problem of old evidence: namely, if scientists really have already priced in the impact of E in their present probabilities for H, then the question of confirmation is irrelevant.

be mediated by principles about the choice of which probability model to adopt, principles which are conspicuous by their absence, or by a plausible but admittedly imprecise appeal to context, as in §4.2. In general, as the discussion of the old evidence problem underscored, our framework separates questions of overall degree of support from incremental confirmation. Both have their place, so it is best if we needn't choose. It is admittedly puzzling that the Likelihoodist doesn't mandate consideration of the base rate, even when robust statistics are available; it may be needlessly complex to regard that aspect of statistical analysis, treated in a uniform way by the Bayesian, as requiring a non-probabilistic individual epistemology to regulate belief given likelihoods. On the other hand, whether base rate neglect is seen as a problem will return us to the motivations for Likelihoodism, namely, the arbitrariness of priors. That will have to await further discussion Section 5.

The mammogram example gave us another intuitive success for Bayesian confirmation theory, and a crisp illustration of differences between it and Likelihoodism. But more decisive objections to Likelihoodism come from counterexamples to their predictions about evidential favouring. The Likelihoodist cares about likelihoods, probabilities of evidence given theories; they do not consider probabilities of theory given evidence. But in some cases such probabilities appear to be independently well-defined, without appeal to ungrounded priors; and they seem to be able to make a difference confirmationally.

Suppose we will draw a card from a well-shuffled standard deck, and adopt a standard probability model for card draws (the example is due to Fitelson (2007), p. 476–477). Let E = 'the card is a spade', H_1 = 'the card is the ace of spades', and H_2 = 'the card is black'. Here the likelihood of E on H_1 is maximal; while on the other hand, H_2 leaves open a real possibility that E is false: $\Pr(E \mid H_2) = 0.5$. The Law of Likelihood tells us that E favours H_1 over H_2. But notice that E tells us, conclusively, that H_2 is correct, and this probability is, given this background probability model, as robust and determinate as the likelihoods. Likewise, the model yields $\Pr(H_1 \mid E) = 1/13$. This is a counterexample to Likelihoodism; E cannot plausibly support a hypothesis it says might be false over one it guarantees is true. (We'll discuss this example further in §4.6.)

The problem here isn't distinctively Likelihoodist, since it also arises for the Bayesian who offers an account of 'favouring' involving comparisons of ratios of likelihoods:

$$\Pr(H_1 \mid E)/\Pr(H_1) = 4 > \Pr(H_2 \mid E)/\Pr(H_2) = 2.$$

But the Likelihoodist, unlike the Bayesian (§4.6), is locked into likelihood comparisons as fundamental to their theory of evidence. So the existence of cases like this, when judgments of evidential favouring provide counterexamples to Likelihoodist predictions, are excellent reason for moving beyond the strictures of Likelihoodism.

4.4 The Scientific Method

That Bayesian confirmation theory can reproduce judgments about evidential support in particular cases is promising. But we also have an existing body of principles and heuristics that together form a proto-theory of evidential support, embedded in the practice of science. The success of science indicates some value to these truisms comprising the 'scientific method', so Bayesian accounts of evidential support are vindicated when they are able to reconstruct and systematise ideal scientific practice. As Earman notes, 'an adequate account of confirmation is not under obligation to give an unqualified endorsement to all such truisms' (Earman 1992: 77), but ideally it should explain the success of those we should endorse. For reasons of space, our discussion here will be partial. The topic of Bayesian philosophy of science is treated extensively elsewhere (Horwich 1982: 100–130; Earman 1992: 63–86; Howson and Urbach 1993: 117–64; Bovens and Hartmann 2003; Schupbach 2022).[14]

Consider the role of *refutation* in scientific inference. When a theory makes a determinate prediction (relative as always to background assumptions) which is not borne out in experiment, that is often taken to decisively undermine the prospects of that theory. For example, the simplest aether theory of the propagation of light was decisively undermined by the outcome of the Michelson-Morley experiment, which did not observe the predicted

[14] It is no longer correct to claim, as Glymour did in -Glymour (1981), that 'There is very little Bayesian literature about the hodgepodge of claims and notions that are usually canonized as scientific method'.

difference in the speed of light in perpendicular directions, while background assumptions excluded rival aether-preserving explanations.[15] This sort of case is central enough to scientific practice that Popper (1959) was able to make it the centrepiece of his 'falsificationist' approach to theory choice (without too much resistance from working scientists, despite falsificationism as a whole looking like a mischaracterization of the aims of science). The correctness of this judgement is supported by a Bayesian model. When H determinately predicts E, relative to background assumptions, that is reflected in the likelihood $\Pr(\neg E \mid H) = 0$ (there is no prospect of E's falsity). In that case,

$$\Pr(H \mid \neg E) = \Pr(\neg E \mid H)\frac{\Pr(H)}{\Pr(\neg E)} = 0.$$

So falsifying evidence conclusively undermines a hypothesis. The displayed equation shows that, in general, the degree of support of a hypothesis by evidence is proportional to increasing likelihood, the limit case being where H entails E.

There is an asymmetry here, which Popper's views may mesh with, in that evidence a theory predicts we will not see yields conclusive disconfirmation, while evidence a theory predicts we will see does not yield conclusive confirmation. (H is conclusively confirmed by E only when E excludes $\neg H$, relative to background knowledge.) But in those cases where a hypothesis predicts some proposition, so that $\Pr(E \mid H) = 1$, we see a degree of support of H by E that is equal to $\frac{\Pr(H)}{\Pr(E)}$. Firstly, note that if we hold the prior probability of H fixed, the more improbable E is, the more support it provides H. This gives us the value of *surprising evidence*: other things being equal, antecedently unexpected evidence has greater evidential impact for the hypotheses that predict it than evidence we'd expect anyway.[16]

15 The role of background assumptions is vital; one can always, as the Quine-Duhem thesis would have it, save a theory by rejecting a background hypothesis. Relative to a fixed background of total evidence against which a theory makes a determinate prediction, it is the theory which is undermined. But very often the evidential background leaves open both hypothesis and needed auxiliary assumptions. In that case the relative confirmation and disconfirmation of theory and auxiliary will depend on the relative impact of the refuting evidence on the posterior probability of each; it is possible to model, in a fairly robust way, historically plausible choices of epistemic perspective that reproduce widely accepted judgments about when theories are refuted and when auxiliaries are to be rejected (Dorling 1979; Strevens 2001), though

16 An example discussed by Jeffrey (2008: §2.3): In 1846 the French astronomer Leverrier, on the basis of various irregularities in Uranus' motion and Newtonian mechanics (call this

Secondly, note that $\Pr(H \mid E) = \frac{\Pr(H)}{\Pr(E)}$ entails $\Pr(H \mid E) > \Pr(H)$: entailed evidence invariably confirms, just as the Entailment condition would have it (§2.1). This prompts us to revisit Hempel's theory, through the lens of the Missing Bushwalker case from §2.2. Intuitions about support in that case posed problems for Hempel's principles of confirmation. Recall in that case the missing bushwalker could be, with equal prospect, in any of six sectors (nw, nc, ne, sw, sc, se). Our hypotheses were that the bushwalker is in the west, i.e., $\mathsf{nw} \vee \mathsf{sw}$, and that he's in the east (i.e., $\mathsf{ne} \vee \mathsf{se}$); our evidence is that $\neg\mathsf{nc}$. First note that the likelihoods $\Pr(\neg\mathsf{nc} \mid \mathsf{w}) = 1 = \Pr(\neg\mathsf{nc} \mid \mathsf{e})$, so the evidence confirms both w and e, as antecedent judgement suggests. But since e is confirmed, $\neg\mathsf{e}$ is not confirmed:

$$\Pr(\neg\mathsf{e} \mid \neg\mathsf{nc}) = \frac{\Pr(\neg\mathsf{nc} \mid \neg\mathsf{e})\Pr(\neg\mathsf{e})}{\Pr(\neg\mathsf{nc})} = \frac{3/4 \cdot 2/3}{5/6} = 0.6.$$

So the Bayesian approach rejects the Consequence and Consistency Conditions. We'll come back to Hempel's Raven's paradox below (§4.5).

It appears to be a methodological rule that, other things being equal, the more *diverse* the sources of evidence for one's theory, the more strongly confirmed that theory is. This can be captured in this maxim: *A theory which makes predictions in a number of disparate and seemingly unconnected areas is more confirmed by that evidence than is a theory which is confirmed by predictions only about a narrow and circumscribed range.* This maxim is also part of the grounds for recommending random sampling in population inference. The Bayesian insight is that diverse evidence is not internally correlated (Howson and Urbach 1993: 160; Steel 1996: 667–68). If, for example, the hypothesis is that all swans are white, then swans collected from different countries would, if white, provide better evidence for the hypothesis than swans collected from the same pond, as we know that if one swan on a pond is

H), predicted the existence, and orbit, of a large, extra-Uranian planet. This planet was was subsequently found (call this E) and named 'Neptune'. The prior credence in E is $\frac{1}{180}$ – the probability of choosing a point on a circle to within 1 degree (since all the planets are found in the ecliptic, and that was the accuracy of Leverrier's prediction). $\Pr(E \mid H) \approx 1$; so $\Pr(H \mid E) \approx \frac{Pr(H)}{Pr(E)} \approx 180 \cdot \Pr(H)$: a strong confirmatory boost for Newtonian mechanics. This kind of pattern is seen also in the Babbage example from §4.1.

white, it is much more likely to be related to other swans in its pond, and those are more likely therefore to be white. If the hypothesis is false, correlations between diverse evidence are more coincidental than correlations between similar evidence. (From a falsificationist perspective, diverse predictions pose a more severe test to the proposal that our hypothesis is false.)

Again focussing on the case where hypothesis predicts evidence with (near) certainty, if the evidence is diverse, it consists of at least two propositions, E_1 and E_2, such that truth of one is not positively relevant to the truth of the other, *if* the hypothesis in question is false. (If it is true, then the evidence is all true, so correlated.) So E_1 and E_2 are diverse *relative to H* iff the likelihood $\Pr(E_1 \wedge E_2 \mid \neg H)$ is low, or at least if it is not greater than the product of the individual likelihoods $\Pr(E_1 \mid \neg H)\Pr(E_2 \mid \neg H)$.

The likelihood ratio $\frac{\Pr(E_1 \wedge E_2 \mid \neg H)}{\Pr(E_1 \wedge E_2 \mid H)}$ features in this formulation of Bayes' theorem:

$$\Pr(H \mid E) = \frac{\Pr(H)}{\Pr(H) + \frac{\Pr(E_1 \wedge E_2 \mid \neg H)}{\Pr(E_1 \wedge E_2 \mid H)}\Pr(\neg H)}.$$

If the hypothesis H predicts both E_1 and E_2, then the likelihood $\Pr(E_1 \wedge E_2 \mid H)$ is close to one. The likelihood ratio is therefore close to $\Pr(E_1 \wedge E_2 \mid \neg H)$. Substitute this in:

$$\Pr(H \mid E_1 \wedge E_2) \approx \frac{\Pr(H)}{\Pr(H) + \Pr(E_1 \wedge E_2 \mid \neg H)\Pr(\neg H)}.$$

But if E_1 and E_2 are diverse (uncorrelated) evidence, then so long as neither is certain given $\neg H$, this guarantees that the term $\Pr(E_1 \wedge E_2 \mid \neg H) < 1$, and hence that $\Pr(H \mid E_1 \wedge E_2) > \Pr(H)$.

Moreover, the more surprising each piece of independent evidence is, and the more we have, the more confirmatory diverse evidence is. As we consider additional pieces of diverse evidence,

$$\lim_{i \to \infty} \Pr(E_1 \wedge E_2 \wedge \ldots E_i \mid \neg H) = 0,$$

hence $\Pr(H \mid E_1 \wedge E_2 \wedge \ldots E_i)$ tends to 1. This result requires independent evidence; correlated evidence doesn't have increasing confirmatory impact the more of it one collects.

We must bear in mind as always that judgments of diversity are relative to theoretical background:

> … the notion of variety of evidence has to be relativized to the background assumptions K, but there is no more than good scientific common sense here, since, for example, before the scientific revolution the motions of the celestial bodies seemed to belong to a different variety than the motions of terrestrial projectiles, whereas after Newton they seem like peas in a pod. (Earman 1992: 79)

Here is another methodological rule: other things being equal, science prefers naturally arising theories to *ad hoc* ones designed to predict the same evidence. Suppose a hypothesis, springing unbidden to the scientific mind, entails a certain piece of evidence; and another hypothesis is then designed to mimic the success of the first theory, entailing the evidence by construction. An example is provided by van Fraassen:

> It is part of [Newton's] theory that there is such a thing as Absolute Space, that absolute motion is motion relative to Absolute Space, …. He offered in addition the *hypotheses* (his term) that the centre of gravity of the solar system is at rest in Absolute Space. But as he himself noted, the appearances would be no different if that centre were in any other state of constant relative motion. This is the case for two reasons: differences between true motions are not changed if we add a constant factor to all velocities; and force is related to changes in motion (accelerations) and to motion directly. (van Fraassen 1980: 46)

Consider Newton's theory N, and the constructed alternative $N + \vec{v}$, that the centre of gravity of the solar system has constant absolute velocity $\vec{v}$. These theories will make the same empirical predictions, so from the point of view of evidence they are indistinguishable. Yet one might think, Newton's theory is clearly to be preferred to each of the arbitrary variants.[17]

[17] One might think an even better theory is a neo-Newtonian theory that does away with absolute space altogether and yet makes the same predictions. Maybe so, but that doesn't alter the fact that Newton's original theory is more supported by the evidence than its rivals.

This sort of case has been raised as an objection to Likelihoodism (cf. Norton 2011: 420–22). Considering only the likelihoods of hypotheses, what resources does the Likelihoodist have to explain our distaste for *ad hoc* hypotheses? N and $N + \vec{v}$ both entail the evidence, so the likelihoods are the same, and there the Likelihoodist account stops. To explain this methodological preference, we cannot appeal to the likelihoods alone, but must also appeal to the disparity in prior probability between the antecedently plausible Newtonian theory N and the antecedently implausible $N + \vec{v}$. One might cite all sorts of explanations for why this prior disparity exists – perhaps Newton's theory is simpler, more natural, less arbitrary – but that it exists and drives judgments of confirmation is undeniable. The Bayesian view of *ad hoc* theories then is that they may have some credibility, and may be supported by evidence, but that generally the fact that they are cooked up to preserve the empirical predictions prompts people to assign them low probability:

> people often respond immediately with incredulity, even derision, on first hearing certain *ad hoc* hypotheses. … it is … likely that they are reacting to what they see as the utter implausibility of the hypothesis. (Howson and Urbach 1993: 158)

There is one kind of case that may trouble the Bayesian: when the *ad hoc* hypothesis is cooked up to be entailed by the original hypothesis. Suppose $N^\dagger$ is stipulated to be the theory, 'N or the empirical appearances are just as if N'. Any evidence E entailed by N is also entailed by $N^\dagger$; since N entails $N^\dagger$, $\Pr(N^\dagger) \geq \Pr(N)$. So we can't appeal to the implausibility of *ad hoc* rivals to explain our decided preference for N; if N is probable enough to be believed, so is $N^\dagger$. In this case, we might need to appeal to another broadly Gricean principle: that our conversational contributions be as informative as they can be, subject to other conversational norms. N is more informative than $N^\dagger$. Perhaps our preference for N is about what we should *say* we believe, more than about what is credible.

4.5 The Ravens Paradox Revisited

The role of background assumptions is also vital for the Bayesian treatment of Hempel's paradox of the ravens (§2.1). The paradox arises for hypothetical-deductive confirmation due to its commitment to the Equivalence Condition, Nicod's Condition, and the judgment that a non-F, non-G doesn't incrementally confirm the hypothesis that *all Fs are Gs*. Clearly the Bayesian will accept the Equivalence Condition; evidential support relates propositions, given an epistemic perspective, and logically equivalent propositions are identical. Existing counterexamples, such as the Yorkshire grasshopper, show that Nicod's Condition isn't invariably correct. But we cannot conclude that instances *never* confirm generalisations; that would be overkill as a response to the paradox. The Bayesian has an account here; indeed, they have two. For there is an account of the circumstances under which Nicod's Condition fails; and there is an account of how to dissolve the paradox, even when Nicod's Condition applies.

How does the Bayesian account for the failures of Nicod's Condition? By appeal to background knowledge. Recall Good's (1967: 322) example from §2.4 of the two worlds, one with relatively few crows, all of which are black, versus the other with many crows, one of which is white. Suppose our background evidence is symmetrical between the two possible worlds, so we adopt an epistemic perspective in which $\Pr(H) = 0.5$. The evidence E is that a black crow is selected at random: $\Pr(E \mid H) = \frac{100}{1{,}000{,}100} \approx 0.0001$, while $\Pr(E \mid \neg H) = \frac{1000}{1{,}001{,}001} \approx 0.001$.[18] Hence $\Pr(E) \approx 0.5 \cdot 0.0001 + 0.5 \cdot 0.001 \approx 0.00055$, and $\Pr(H \mid E) \approx 0.009 < \Pr(H)$. So we see evidence that disconfirms the hypothesis even while being an instance of it. This example is rather confected, but more realistic examples with the same structure are available. Suppose a epidemiologist is thinking about a virulent illness, endemic overseas, which they are worrying may have begun to take root in their community. The scenarios they are considering are two: there is no reservoir of disease in their community, and there is a significant hitherto-undiagnosed population of positive cases. They get word of a positive case recently diagnosed, an instance of the generalisation 'all cases of the disease have been identified'. But of course this positive case is conclusive evidence that the disease is in the community, and hence conclusive evidence against the truth of the generalization.

The counterexamples to Nicod's condition have the distinctive feature

[18] Here notice we must use the total evidence acquired; the evidence is not merely that there is a black raven, but that a black raven was the result of a random selection – it is clearly that latter aspect of the evidence that renders it so unlikely in the world where the generalization is true.

that encountering certain instances of the generalization indicates its falsity, relative to background knowledge. In many cases, however, whether we encounter an F is independent of any generalization about the qualities of Fs. Consider such a case (Howson and Urbach 1993: 127; see also Hosiasson-Lindenbaum 1940; Earman 1992: 72). Let the proposition that an entity a is encountered be part of the background knowledge, and then let R be the proposition 'a is a raven', and B the proposition 'a is black'. That a black raven is encountered is then $R \wedge B$, a non-black, non-raven encountered being $\neg B \wedge \neg R$. Let A be the hypothesis 'all ravens are black'. $R \wedge B$ confirms A iff $\Pr(A \mid R \wedge B) / \Pr(A) > 1$.

Rearranging an instance of Bayes' theorem gives us

$$\frac{\Pr(A \mid R \wedge B)}{\Pr(A)} = \frac{\Pr(R \wedge B \mid A)}{\Pr(R \wedge B)}.$$

If all ravens are black, then a's being a raven guarantees it to be black: so $\Pr(B \mid A \wedge R) = 1$. Hence $\Pr(R \wedge B \mid A) = \Pr(R \mid A)$; and the independence of encountering a raven from the hypotheses about the characteristics of ravens entails that $\Pr(R \mid A) = \Pr(R)$. So

$$\frac{\Pr(A \mid R \wedge B)}{\Pr(A)} = \frac{\Pr(R)}{\Pr(R \wedge B)}.$$

Similar reasoning will show that

$$\frac{\Pr(A \mid \neg B \wedge \neg R)}{\Pr(A)} = \frac{\Pr(\neg B)}{\Pr(\neg B \wedge \neg R)}.$$

Turning our attention to the probability of the evidence, suppose we have various hypotheses about the proportion of ravens that are black. Let F_i state that the frequency of black things among the ravens is $100i\%$; thus $A = F_1$. The probability, given one has encountered something, that it is a black raven is the probability of encountering a raven, multiplied by the probability that the raven is a black one: $\Pr(R)\Pr(B \mid R)$. Given we don't know the frequency of black ravens among ravens, we use our background distribution over F_i to calculate $\Pr(B \mid R) = \sum_i \Pr(B \mid F_i \wedge R)\Pr(F_i \mid R)$. In normal cases, R is admissible evidence for the frequency hypothesis F_i. So $\Pr(F_i \mid R) = \Pr(F_i)$. And $\Pr(B \mid F_i \wedge R) = i$; this is an instance of the Principal Principle (§3.9).

Putting that all together:

$$\frac{\Pr(A \mid R \wedge B)}{\Pr(A)} = \frac{\Pr(R)}{\Pr(R)\Pr(B \mid R)} = \frac{1}{\Pr(B \mid R)} = \frac{1}{\sum_i i \Pr(F_i)}.$$

Similar reasoning gets us to this:

$$\frac{\Pr(A \mid \neg B \wedge \neg R)}{\Pr(A)} = \frac{\Pr(\neg B)}{\Pr(\neg B)\Pr(\neg R \mid \neg B)} = \frac{1}{\Pr(\neg R \mid \neg B)}.$$

But in this case we don't need to consider chance hypotheses about the frequency of non-ravens among the non-black things, because our background evidence includes that the number of non-black things is vastly more than the number of ravens, so almost all non-black things aren't ravens: $\Pr(\neg R \mid \neg B) = 1 - \epsilon$, hence $\Pr(A \mid \neg B \wedge \neg R) \gtrapprox \Pr(A)$. So we might get a tiny improvement in the degree of support for A given a non-black, non-raven over the unaugmented background information.

The same is not true for the observation of a black raven. There the improvement of the prospects of A depends on the distribution over the chance hypotheses F_i. Suppose we have a rough model, assigning equal probability of 0.25 to each of $F_0, F_{1/3}, F_{2/3}, F_1$. Then $\sum_i i \Pr(F_i) = 0.5$, and hence $\Pr(A \mid R \wedge B)$ is significantly greater than $\Pr(A)$. (This is representative for any epistemic perspective that assigns a uniform prior to each of the hypotheses about frequency.) That is because we antecedently gave significant credence to hypotheses stating the proportion of black ravens among the ravens is low, and an encounter with a black raven was significantly in tension with those hypotheses. On the other hand, had background knowledge already indicated the proportion of black ravens was high, the confirmatory impact of the evidence would have been less.

4.6 Measuring Confirmation

Howson and Urbach (1993) summarise as follows:

> the fact that $R \wedge B$ and $\neg B \wedge \neg R$ both confirm a hypothesis does not imply that they do so with equal force. Once it is recognised that confirmation is a matter of degree, the conclusion [of Hempel's paradox] is no longer so counterintuitive, because it is compatible

> with $\neg B \wedge \neg R$ confirming 'All Rs are Bs', but to a minuscule and negligible degree. (Howson and Urbach 1993: 127)

Here they suggest that part of the explanation for the judgments in the ravens paradox is a confusion between *no* confirmation and *negligible* confirmation. But that explanation invokes a notion of *degree of confirmation* that is as yet unanalysed. Degree of support of a hypothesis, relative to background evidence, we have an analysis of – that is just $\Pr(H)$ – but degree of confirmation is a distinct notion, a measure of how much incremental confirmation E provides to H over the background evidence.[19]

Howson and Urbach note that when $\Pr(H \mid E) \approx \Pr(H)$ there is little confirmation of H by E, while when $\Pr(H \mid E) \gg \Pr(H)$ there is significant confirmation. This is most naturally understood by appeal to the Ratio account of degree of confirmation (§4.3). On that view, E confirms H to the degree to which the ratio of posterior to prior exceeds 1. Often the logarithm of this ratio is taken, giving us:

Log-Ratio The log-ratio *measure* is defined as $r(H,E) \stackrel{\text{def}}{=} \log\left(\frac{\Pr(H|E)}{\Pr(H)}\right)$.[20]
The Log-Ratio *analysis* says that, where $\mathfrak{c}(H,E)$ measures the degree of confirmation, $\mathfrak{c}(H,E) = r(H,E)$ (Horwich 1982: 57; Milne 1996; Eells and Fitelson 2002: 131).

The Log-Ratio analysis of evidential favouring entails Contrastive Favouring (Ratio), encountered in §4.3, and hence entails the Law of likelihood. Accordingly, in Fitelson's (2007: 476–77) 'ace of spades' example (§4.3), the log-ratio measure entails that $\mathfrak{c}(\text{ace of spades}, \text{spade}) > \mathfrak{c}(\text{black}, \text{spade})$, yet it would not be plausible to say that the proposition that a card is a spade is both more conclusive evidence that the card is black than that it is the ace of spades, while also being more negligible in its support. In general, the follow-

[19] This is why direct comparison of degrees of support – e.g., the proposal that E favours H over H' if and only if $\Pr(H \mid E) > Pr(H' \mid E)$ – is not a good measure of confirmation. That comparison is also about background evidence and may not be represent the incremental confirmation contributed by E. Again the difference between degree of support and degree of confirmation is pertinent. We can accept this principle however: that if $\Pr(H \mid E) > \Pr(H' \mid E)$ then the *total evidence*, including E, favours H over H'.

[20] If $\Pr(H \mid E) > \Pr(H)$, i.e., E confirms H, $r(H,E)$ is positive; if $\Pr(H \mid E) < \Pr(H)$, $r(H,E)$ is negative; and $r(H,E) = 0$ when H and E are independent. Since $x > y$ iff $\log(x) > \log(y)$, this won't change any rankings of comparative confirmation.

ing result holds (Eells and Fitelson 2002: 139):

$$\begin{aligned} r(H,E) &= \log\left(\frac{\Pr(H \mid E)}{\Pr(H)}\right) \\ &= \log\log\left(\frac{\Pr(E \mid H)}{\Pr(E)}\right) \\ &= r(E,H). \end{aligned}$$

But, as the cards example shows, the degree to which E supports H is not always the degree to which H to supports E – that a card is a spade appears to supports its being black conclusively, while it's being black doesn't conclusively support its being a spade. The Log-Ratio measure is symmetric; it ought not be, i.e., in general, $\mathfrak{c}(H,E) \neq \mathfrak{c}(E,H)$.

If other measures of confirmation gave equally poor verdicts, in this case or others, then perhaps we could learn to live with this counterexample. But we do not need to. There are a number of other measures that have been proposed. We will consider the two most prominent after the Log-Ratio measure.

Difference The difference *measure* is defined as $d(H,E) \stackrel{\text{def}}{=} \Pr(H \mid E) - \Pr(H)$.
The Difference *analysis* says that $\mathfrak{c}(H,E) = d(H,E)$. (Earman 1992: 64; Jeffrey 1992: 72; Eells and Fitelson 2002: 131)

Log-Likelihood The log-likelihood *measure* is defined as $l(H,E) \stackrel{\text{def}}{=} \log\left(\frac{\Pr(E|H)}{\Pr(E|\neg H)}\right)$.
The Log-Likelihood *analysis* says that $\mathfrak{c}(H,E) = l(H,E)$. (Fitelson 1999, 2007; Eells and Fitelson 2002: 131).

These measures both are, like r, positive when E confirms H, zero in cases of independence, and negative in cases of E disconfirming H. Note that neither of these measures is symmetric (Eells and Fitelson 2002: 138, 140); they have some prospect, unlike r, of being an analysis of $\mathfrak{c}$.

Similar examples to those which spelled trouble for the ratio measure also trouble the difference measure, however. Suppose the evidence is that a heart is drawn from a well-shuffled deck, and we are considering the hypotheses that the card drawn is *not a club*, and that it is *a heart but not a face card*.

Here are the relevant probabilities:

$$\begin{aligned} \Pr(\text{not-club}) &= 3/4 & \Pr(\text{non-face-heart}) &= 10/52 \\ \Pr(\text{not-club} \mid \text{heart}) &= 1 & \Pr(\text{non-face-heart} \mid \text{heart}) &= 10/13 \\ d(\text{not-club}, \text{heart}) &= 1/4 & d(\text{non-face-heart}, \text{heart}) &= 30/52 \end{aligned}$$

Here it is clear that while the evidence of a heart card conclusively entails the hypothesis that it is not a club, the high prior probability limits the degree of confirmation. That same evidence also supports the hypothesis that it is a numbered heart, though not conclusively. Here, though, the low prior probability allows a much greater boost in absolute support, which is reflected in the difference measure. Intuitively, this is the wrong result: how can an inconclusive piece of evidence support a hypothesis to a greater degree than a conclusive piece of evidence?

The Log-Likelihood involves these probabilities:

$$\begin{aligned} \Pr(\text{heart} \mid \text{not-club}) &= 1/3 & \Pr(\text{heart} \mid \text{non-face-heart}) &= 1 \\ \Pr(\text{heart} \mid \text{club}) &< \epsilon & \Pr(\text{heart} \mid \neg\text{non-face-heart}) &= 3/42 \\ l(\text{not-club}, \text{heart}) &> \log(\frac{1}{3\epsilon}) & d(\text{non-face-heart}, \text{heart}) &= \log(42/3) \approx 2.64 \end{aligned}$$

Here we fudge slightly to avoid division by zero; we cannot assign a single degree of confirmation to the case of conclusive evidence for a hypothesis, but we do know that as ϵ tends to zero, the degree of confirmation tends to infinity, i.e., maximal degree of confirmation.[21] So l accords better with antecedent judgment here than d.

The debate, as over any 'conceptual analysis', rumbles on (Crupi 2021). In this case, with no unique best satisfier of 'intuitive' desiderata, a pluralist attitude might suggest itself. All measures agree on the qualitative fact of *whether* E confirms H, differing only on the question *how much?*. But what turns on this question? What ultimately matters for belief and action is how much the total evidence from some epistemic perspective we have adopted supports a hypothesis. Incremental confirmation matters because a confirmed theory will be more overall supported by the evidence, and the successes

[21] For this reason another measure due to Kemeny and Oppenheim (1952) is sometimes suggested that delivers the same ranking as l but is well-defined in such cases.

of the Bayesian in accounting for scientific maxims depend on that notion. But it is harder to identify scientific maxims that require a very precisely specified measure of confirmation. In most concrete cases, the existence of some plausible measure that delivers an acceptable verdict is taken as sufficient to vindicate a Bayesian approach. The plurality of measures would then reflect the plurality of our interests in quantifying confirmation; e.g., sometimes we care about the absolute size of increases, in which case d is a useful measure; sometimes about relative size, in which case l might do better.[22] This isn't like the Church-Turing thesis, where radically different attempts to characterise a pre-theoretical notion of computability ended up converging on the same class of computable functions. That example is highly unrepresentative of the process of explication. Most philosophically interesting notions turn out upon precisification to splinter into finely distinguished but broadly overlapping notions,[23] and it is hardly to be suspected that measures of incremental confirmation will be different.

In any case, our motivation was to try and model the suggestion that while a white shoe might confirm the hypothesis that all ravens are black to some extent, it won't be as much as a black raven would. Our discussion in §4.5 gave us that $\Pr(A \mid \neg B \wedge \neg R) \gtrapprox \Pr(A)$, so on the Difference measure we get immediately that $d(A, (\neg B \wedge \neg R)) \gtrapprox 0$. Our toy example in that same discussion gave us that $\Pr(A \mid R \wedge B) = 2\Pr(A)$, so that $d(A, R \wedge B) > \Pr(A)$. Given some non-negligible prior probability for A, this will entail that $d(A, R \wedge B) > d(A, (\neg B \wedge \neg R))$.

A similar exercise for the Log-Likelihood measure gives the same verdict. The probability calculus gives us:

$$\Pr(\neg B \wedge \neg R \mid A) = \Pr(\neg R|\neg B \wedge A)\Pr(\neg B \mid A) = \Pr(\neg B \mid A) \approx \Pr(\neg B).$$

22 Some have argued against d on these grounds, because a large increase in relative risk; e.g., a thousand-fold increase in cancer risk after radiation exposure) might be associated with a very low value of $d(\text{cancer}, \text{radiation})$ (Schlesinger 1995; Zalabardo 2009). It is actually by no means clear that the pre-theoretical judgments about confirmation we are attempting to systematize are at variance with this result.

23 Access consciousness and phenomenal consciousness; control versus ultimate/proximate sourcehood notions of free will; etc.

But if A is false, there is a tiny chance of encountering a raven given one encountered a non-black thing, so $\Pr(\neg R \mid \neg B \wedge \neg A) > 0$. Hence $\Pr(\neg B \wedge \neg R \mid \neg A) < \Pr(\neg B \mid \neg A) \approx \Pr(\neg B)$. Hence $l(A, (\neg B \wedge \neg R)) \gtrapprox 0$.

Repeating the exercise for the evidence of a black raven gives:

$$\Pr(R \wedge B \mid A) = \Pr(B|R \wedge A)\Pr(R \mid A) = \Pr(R \mid A) \approx \Pr(R).$$

But now $\Pr(B \mid R \wedge \neg A)$ isn't necessarily close to 1; it might be quite significant, depending on what hypotheses about the proportion of ravens that are black remain live and what their priors are. So given $\Pr(R \mid \neg A) \approx \Pr(R)$,

$$\Pr(R \wedge B \mid \neg A) = \Pr(B|R \wedge \neg A)\Pr(R \mid \neg A) \ll \Pr(R).$$

Hence $l(A, R \wedge B)$ can be significantly different from zero (again, depending on how antecedently supported the hypothesis that all ravens are black was), and so $l(A, R \wedge B) > l(A, (\neg B \wedge \neg R))$.[24]

4.7 Inductive Logic and Inductive Framework

The best place to finish our positive Bayesian story is where we began: with induction. Inductive inference was understood as covering all species of inference to the best explanation, including inverse inference from a sample to a population, or to a subsequent sample (§1.2). The synchronic aspect of this, the part that could be the subject of inductive logic, is to articulate constraints that explanation places on rational epistemic perspectives. (We emphasise the 'best explanation' part of 'inference to the best explanation'.) Induction is vindicated to the extent that a body of evidence supports the best explanation of that evidence. In Bayesian terms, broadly speaking, we accommodate induction by showing that when H explains E, that E confirms H; and that when H is the *best* explanation of E, H is probable in light of a body of total evidence including E.

One standard view of explanation is that an explanation shows how an otherwise puzzling event is to be accommodated and made comprehensible within a broader framework. Van Fraassen (1980: ch. 6) suggests that explan-

[24] A full Bayesian treatment of this paradox, exploring just how weak the needed assumptions can be, is given by Fitelson and Hawthorne (2010).

ations are proffered as answers to 'why'-questions; to explain an event is to provide relevant information about the origins of or participants in event, relative to the background presuppositions of the questioner. Very often, though perhaps not invariably, this will take the form of 'information about its causal history' (Lewis 1986b: 217). So to explain why the vase broke could involve citing a cause of the vase breaking, such as its being dropped. Yet a question coming from a different background might demand a different answer. Suppose our request for explanation takes this form: 'Yes, I know it was dropped, but why did it *break*?'); in that case the request for explanation might be satisfied by providing information about the fragility of the vase. This is still information about the causal history.

Very often, causal relations are manifest in relations of statistical dependence. When C is a body of information about causes, and E some effect of those causes, very often $\Pr(E \mid C) > \Pr(E)$.[25] The broken vase is more probable given it was dropped than otherwise. The background is involved in selecting a pertinent probability function: it is more probable that the vase breaks given it was fragile and dropped, than that it is dropped alone, given a background that does not build in the fragility of the vase. Putting this together: very often, to explain an event is to offer information, relative to a background body of evidence, such that the likelihood of the event given the information is greater than its prior probability:

> where the hypotheses are specific, a hypothesis, H, explains the data better than H', if true, just when H would make the data more expected than H'. In judging which hypothesis renders the data most understandable, we consider nothing more than which hypothesis renders it most expected. (Henderson 2014: 700)

A theory is explanatory to the extent that it encapsulates such information, so that – very often – H explains E just when it renders E more likely than otherwise, $\Pr(E \mid H) > \Pr(E)$. An elementary application of Bayes' theorem then entails:

Bayesian Explanation When H explains E by making it more likely than otherwise, relative to some background evidence and conception of

[25] Causation is not quite perfectly manifest in statistical dependence, since there may be causes that do not raise the probability of their effects (Rosen 1978; Glynn 2010: 349–53).

evidential relevance, then E confirms H: if $\Pr(E \mid H) > \Pr(E)$, then $\Pr(H \mid E) > \Pr(H)$.

Recall an example from §1.2. The Channeled Scablands of western Washington state is a complex landscape of braided channels, exhibiting the characteristics seen in microcosm in dry gorges incised into harder rock, such as potholes, gravel bars, and scoured deep grooves. Bretz hypothesised that this landscape was indeed the effect of a cataclysmic flood, in which debris-laden water was discharged on such a vast scale that an existing dissected plateau was filled beyond the capacity of its existing drainage, so that water spilled over the top of the plateau, erased the existing topsoil, and carved branching and reuniting channels into the bedrock (Baker 2009: 402–3). The hypothesis certainly makes the evidence more probable than otherwise: given attested geological mechanisms as the background assumptions, a gigantic flood would produce just what is seen at landscape scale.

Nevertheless, Bretz' hypothesis took many years to gain acceptance, despite its explanatory merits, because the source of such an extraordinary volume of water was unknown. Subsequent evidence of an ice sheet intruding into Idaho suggested the existence of a lake formed behind an ice dam, containing 2100 km^3 of water and covering much of western Montana – it is the collapse of the ice dam and the subsequent evacuation of the whole of this dam in an estimated 48 hours, carving out more than 210 km^3 of soil and leaving behind the Channeled Scablands. It was only after the background evidence provided a remotely plausible source for the required water that Bretz' hypothesis was adopted. So part of what made it ultimately the best explanation was not only its explanation of the data, but its prior plausibility. As before, likelihoods alone do not suffice. To generalize

Bayesian IBE H is the best explanation of the data E when (i) H explains E, $\Pr(H \mid E) > \Pr(E)$, and (ii) H is most probable among competing explanations.

Bayesian inference to the best explanation also shows the limits of IBE. For it is quite possible for H to be the best explanation of E and for some rival to be far more credible, antecedently. A theory may exhibit many explanatory virtues, such as simplicity, elegance, deployment of familiar mechanisms that enable it to generate understanding, etc.[26], and yet not be probable: 'explanatory goodness, whatever it is, looks to be at least somewhat independent of prior conditional [probability]' (Weisberg 2009: 130). Moreover, it may be quite rational to have attitudes that mirror this perspective. In the absence of information favouring the enormous glacial lake, most geologists were sure that (i) Bretz' catastrophic flood hypothesis was an excellent explanation of the data, and (ii) that less unified, ragged, and unfamiliar hypotheses were to be preferred.[27]

What we want, of course, is a story that explains the rationality of both parties. Bretz was open to the existence of a gigantic flood, compelled by the field evidence that seemed to demand it. For him, we may suppose, what was vivid is just how low $\Pr(E \mid \neg H)$ was, which ensured that relative to his background evidence, $\Pr(H \mid E)$ was sufficiently high for overall credibility. The rival view involved a different perspective, $\Pr'$, such that $\Pr'(H)$ was so low that $\Pr'(H \mid E)$ could still not suffice for credibility. This difference in background perspective might accommodate the joint rationality of everyone involved at the earlier stages, though perhaps even $\Pr'(H \mid E \wedge L)$ should be high, where L is the evidence favouring the glacial lake. If 'inference to the best explanation' is understood to involve inference based on explanatory considerations divorced from prior credibility and background knowledge, then it is lucky that Bayesian cannot reconstruct IBE in that sense.[28] But very often, explanatory factors are correlated with the likelihoods of evidence given hypotheses (Henderson 2014: 709), and thus the Bayesian offers an explication of the merits of IBE, when it has merit.

The Bayesian reconstruction of inference to the best explanation is a key part of the Bayesian account of induction, alongside the more particular max-

[26] All those features contributing to what Lipton (2004: 59) calls 'loveliness'.

[27] Largely because these alternatives were taken to be consistent with Uniformitarianism while a catastrophic flood was *exactly* the kind of diluvial hypotheses Uniformitarianism was taken to exclude – though, of course, nothing Bretz posited was unattested, except in scale.

[28] Van Fraassen offers another Bayesian argument against IBE (1989: 166), construed as an inference that boosts the posterior credibility of explanatory hypotheses over and above the extent to which the evidence favours them; so-construed, what we've been discussing is not IBE, and luckily so, since the rule van Fraassen is discussing seems manifestly irrational if truth is what is sought.

ims of inductive methodology discussed earlier (§§4.1, 4.3, 4.4, 4.5). The role of background assumptions has been a constant refrain. And this is just as it should be. The upshot of our earlier discussion of Hume's problem of induction in §2.4 was that we needed to give a theory of the relation 'E supports H relative to standards S', of which we've given a Bayesian account.

We can sharpen this point. Let's consider a highly abstracted but quite general representation of 'classic' inductive inferences. A possible world is an infinite binary sequence of outcomes (like the results of successive coin tosses); the correct theory of a given world is simply identified with the theory that predicts each outcome (it needn't have more 'abstract' theoretical structure); hence the space of possibilities is given by the set of all such sequences. A broadly Bayesian theory of inductive evidential support assigns probabilities to hypotheses based on initial subsequences; it is a classic case of inverse inference from a sample to a population. Some Bayesian results can be established that seem to vindicate induction. Suppose H_i is the sequence that H predicts for the initial i outcomes. Thus $H = H_\infty$, and for each i, $H \vDash H_i$. It can be shown that, so long as $\Pr(H) > 0$,

$$\lim_{i\to\infty} \Pr(H_{i+1} \mid H_i) = 1.$$

That is, the probability of the correct hypothesis tends to 1 as more outcomes conformable with it accumulate (Howson 2000: 72). This looks like a substantive vindication of ampliative inference.

The first limitation to note is that the requirement that $\Pr(H) > 0$, which may look innocuous, is extremely substantive. This space of hypotheses is the Cantor space, the set of all infinite binary sequences; that space is uncountable. If each hypothesis were given equal probability, as under the standard Lebesgue measure, each hypothesis would have prior probability zero, and the result would apply to none of them. If some hypothesis is eventually to be maximally supported by the total evidence, it must be assigned some initial positive probability – indeed, there must be an at most countably many hypotheses assigned positive probability.[29] So almost all possible hypotheses about the sequence of outcomes need to be excluded *ab initio*. Of course one

[29] While one can assign positive (real-valued) probability to each member of a countable set of mutually exclusive hypotheses, one cannot do so for an uncountable set.

might unlucky enough to assign probability 0 to the true hypothesis, in which case after finitely many data points all the live hypotheses will have been refuted. At that point one must simply restart with a new hypotheses space, the set of all infinite binary sequences which begin with the previously observed data. That is still an uncountable space (it simply involves pre-pending the observed data to each element of the Cantor space), and subject to the same worries. So the choice of epistemic perspective already has to make substantive assumptions about which possible hypotheses to consider 'live'; assumptions which are required before confirmation can occur, and even in the presence of observed data, are not fixed by that data.

Secondly, the result tells us nothing about the speed of convergence. Eventually, every rival hypotheses is eliminated by some data point. But after any finite time, the data points eliminating incorrect hypotheses may be arbitrarily far away. So to ensure robust inductive support of the correct hypothesis, we shall have to make the further substantive assumption that the data we have so far are a *representative sample* of the whole population. That assumption seems *a priori* quite strong; an infinite population in which after a certain point no Fs are G can nevertheless begin with arbitrarily many initial Fs which are Gs; hypotheses of that sort simply have to be excluded by fiat.[30] Indeed, whatever we do, we shall need to make some assumptions about what sort of overall hypotheses are supported by an given initial sequence of the data. For example, how quickly should we 'learn from experience'? How many 1s in a row in the initial data should it take us to become more confident than not that all the outcomes will be 1s? How inclined we are to judge that the temporally initial conditions might well be *un*representative of the whole sequence of outcomes – noting for example that in the actual world, the early universe is very unlike the universe over most of time, we might be hesitant to draw any conclusions from the early data — that of course requires some judgment about when the data stops being early.[31] These kinds of assumptions

30 For Hume (1777: ¶4.19), the principle of the uniformity of nature is supported by 'probable arguments', which themselves rest on some sort of representativeness-of-the-past assumption; otherwise it would be quite possible for premises to have been accompanied by conclusions with high frequency and to cease to be associated at all henceforth.

31 Compare also hypotheses about pandemic spread: we should expect those who get the disease early to be systematically different from more cautious individuals who delay infection, in a way that cannot be judged without assumptions about the relative proportions of these individuals in the population.

are most readily understood as constraints on conditional probabilities: what distribution over hypotheses does a given piece of evidence license? These examples show that the Bayesian not only invokes prior probabilities, but may well invoke conditional probabilities too that need not not by uniquely constrained.[32]

Any actionable inductive practice must unavoidably involve some prior assumptions; a third relatum of the evidential support relation, an epistemic perspective encoding both antecedent judgments of hypothesis plausibility *and* prior conditional judgments of evidential relevance. A virtue of the Bayesian account I've developed is that it makes these assumptions explicit in the evidential probability model of epistemic perspectives.

A challenge sometimes posed is that this explicit invocation of a probability model shows that we don't really have a theory of evidential support after all – that we have only an 'inductive framework', rather than an 'inductive logic' that guides scientific argument:

> particular inferences can almost always be brought into accord with the Bayesian scheme by assigning degrees of belief more or less *ad hoc*, but we learn nothing from this agreement. What we want is an explanation of scientific argument; what the Bayesians give us is a theory of learning, indeed a theory of personal learning. But arguments are more or Jess impersonal; I make an argument to persuade anyone informed of the premisses, and in doing so I am not reporting any bit of autobiography. ... Alternatively, and more hopefully, Bayesians may suggest that we give arguments exactly because there are general principles restricting belief, principles that are widely subscribed to, and in giving arguments we are attempting to show that, supposing our audience has certain beliefs, they must in view of these principles have other beliefs,

[32] Of course Bayes theorem shows these are not independent of one another; especially in this case where all the likelihoods are trivial, the conditional probabilities are fixed by the prior probabilities. But all that shows is that, even given a prior distribution over hypotheses, the dispositions to respond to evidence in certain ways will fix the prior probabilities of evidence, rather than taking the latter to somehow be given antecedently.

> those we are trying to establish. There is nothing controversial about this suggestion, and I endorse it. What is controversial is that the general principles required for argument can best be understood as conditions restricting prior probabilities in a Bayesian framework. Sometimes they can, perhaps, but I think that when arguments turn on relating evidence to theory, it is very difficult to explicate them in a plausible way within the Bayesian framework. (Glymour 1981: 74–75; see also Strevens 2004)

The first point to make in response is that this objection seems to require too much of *logic*, regardless of induction (§1.1). The second is that only some Bayesians offer theories of personal learning; not this one. The final point is to note that the proof of Bayesian principles is in the pudding; the review of cases we have undertaken, in which Bayesian precepts rationalise and systematize scientific conceptions of good evidence, provide defeasible grounds favouring the Bayesian model.

Section 5

Uniqueness and the Problem of the Priors

In this section, we begin in §5.1 by outlining various issues around the justification of prior probabilities, and frame responses as permissivist or impermissivist. In §5.2 I describe some permissivist attempts to explain away the demand for unique rational priors. In §5.3 I look at the Principle of Indifference and its role in attempts at constructing a unique prior, and describe the charge of inconsistency levelled at it. I turn to formal approaches to constructing prior probabilities at the end of the section; to Carnap's inductive logic in §5.4, and Solomonoff's algorithmic probability approach in the concluding §5.5. Neither ultimately fares well.

5.1 The Problem of the Priors

The 'problem of the priors' is not one problem, but rather a cluster of issues that circle around the plurality of coherent evidential probability functions.

1. One issue concerns belief and action: if there are many epistemic perspectives, but we need to plump for a particular credence function to feed into our deliberations, how ought we choose an epistemic perspective to adopt? This issue was introduced in §3.9, but any solution will depend on what we say in this section.
2. A second issue concerns the *rationality* of epistemic perspectives. We might want to say that conspiracy theorists, cranks, and those who per-

sist in salvaging a preferred hypotheses by denial of auxiliaries assumptions are being *unreasonable*, even if, technically, they seem to accurately deploy scientific standards. The source of such recalcitrance, on the Bayesian view, lies in the prior distribution over hypotheses. If epistemic perspectives provide ideals for rational belief, the space of epistemic perspectives must be more tightly constrained than hitherto.

3. A related issue concerns *procedural* rationality. The vindication in Section 4 of induction, or of the canons of scientific methodology, required assumptions about the priors. Other priors would vindicate different methodological maxims: counter-induction, preferences for unrepresentative samples and biased evidence. But it would be unreasonable to use these alternative maxims; a reasonable person wouldn't respond to evidence in the ways these perspectives appear to license. The Bayesian picture accommodates rational responses to evidence, but seems unduly tolerant of other responses, to the extent where scientists 'may disagree on sufficiently many important questions that the consensus required for scientific progress is undermined' (Strevens 2006: 82).
4. A further issue concerns the *objectivity* of science. Science is a self-regulating community, with broad intersubjective agreement on procedures and on the space of legitimate theorizing. 'One's expectation, or hope if you will, is that the explanation of the intersubjective agreement on such matters is not merely historical or sociological but has a justificatory character' (Earman 1992: 137–38). Without robust internal constraints on the allowable epistemic perspectives, scientific consensus looks more like the product of exclusion than the ineluctable workings of the scientific method.
5. The Bayesian picture seems to put the cart before the horse. It puts attitudes, whether actual or idealized, in the position that should rightly be occupied by the evidential connections that justify those attitudes: 'our judgment of the relevance of evidence to theory depends on the perception of a structural connection between the two … degree of belief is, at best, epiphenomenal' (Glymour 1981: 92–93).
6. Evidential probabilities must represent ignorance, to be sufficiently amenable to updating in light of new evidence: 'our initial beliefs should not unfairly favor one empirical hypothesis over another. … an

> adequate account of how to respond to evidence should be neutral and "let the data speak for itself"' (Meacham 2014: 1193–94). But Bayesian priors 'exercise a controlling influence' over subsequent attitudes, and are insufficiently neutral (Norton 2011: 428–29).

There are many idea swirling around here, and they are not all pulling in the same direction. A central tension is whether the priors should be neutral, guided by evidence, or instead impose rigid confines on the acceptable responses to evidence. To ensure the rationality and objectivity of inductive practice, we want to require epistemic perspectives to exhibit a uniformity of response to a given piece of information, at least if they share their other background information. That desired uniformity of response across perspectives mandates a non-uniform response to hypotheses by those perspectives, because some hypotheses will gain significant support from the evidence they predict only if 'unreasonable' responses to evidence are deployed. There is nothing incoherent about the theory Hume entertains, that while bread so far has given us nourishment and support, from now on it will not. But in a scenario where it is true, belief in it on the basis of the evidence can only come from a quite different theory of evidential support than they one we actually utilise. This is perhaps another manifestation of the phenomenon of the underdetermination of theory by evidence.

The problem of the priors, as I see it, it is to resolve this tension between treating theories fairly, not letting prejudice scupper their chances at confirmation, and responding to evidence in a productive way that eventually leads to reasonable scientific consensus. Two approaches suggest themselves. The *permissive* response acknowledges that there are many acceptable responses to a given body of total evidence, and no guarantee that the true hypotheses will be eventually favoured by the evidence regardless of which possible epistemic standards are considered (§3.9). Each standard favours some theories over others, but no one standard is singled out prior to experience, hence no theories are disqualified *ab initio*. To secure scientific progress, the permissivist allows (i) it is rational to opt for one epistemic perspective over another, even holding fixed total evidence, and even without its approach being favoured by some decisive epistemic reason; and (ii) shared situational and sociological factors encourage different scientists to opt for more or less similar perspectives. These two factors explain both the widespread agreement on evidential standards, and the rationality of those standards; the drawback many see is that the convergence on common standards isn't explained by their rationality, which seems to leave the approach open to a charge of arbitrariness.

The *impermissive* response is different. Impermissivists deny that there are alternative equally good ways of responding to evidence. We will focus on the variety of impermissivist who asserts that there is a unique acceptable epistemic perspective for any given body of total evidence. Such an impermissivist accepts:

Uniqueness 'There is a unique rational response to any particular body of evidence' (Kopec and Titelbaum 2016: 189); for any 'evidential situation … there is a uniquely rational state to be in right then' (Greco and Hedden 2016: 392).

Given Uniqueness and our previous discussions (§4.7), it will turn out that many hypotheses are guaranteed not to be supported by the evidence, having been excluded from the start by the unique epistemic perspective compatible with null evidence (Meacham 2014: 1213). So not every hypothesis is treated fairly, and (depending on the interaction of Uniqueness with modality) it could turn out that in some scenarios the truth cannot be rationally supported by the evidence. The defenders of Uniqueness are sensitive to this concern, and the concrete implementations of Uniqueness that have been put forward, and that we will discuss, all attempt to build in neutrality between possible hypotheses as a desideratum. The objectivity and rationality of science is secured, as a more than sociological matter, so long as the scientific method follows the dictates of the uniquely rational epistemic perspective – but again, defenders of Uniqueness have used conformity with standard scientific maxims as constraints on the construction of the unique function. Uniqueness may hold out the promise of resolving the problems of the priors.

The principal problems for Uniqueness are two:[1] the manifest implausibility of denying that there can ever be reasonable disagreements about

[1] Permissivism and impermissivism, in my usage, are theses about epistemic perspectives, not individual attitudes. Kopec and Titelbaum (2016: 190–92) note that 'Uniqueness' has been used to label many different claims. This creates an opportunity to deflect certain challenges. For example, perhaps the uniqueness of ideal rationality is compatible with permissivism about individual credence – maybe you can be rational if your credence suitably approximates the ideal, subject to your cognitive limitations. Perhaps permissiveness

the significance of a piece of evidence (Rosen 2001: 71); and the challenge of constructing or defining the uniquely rational epistemic perspective. Some defenders of uniqueness have wanted to dodge the second challenge, suggesting that while there is a uniquely rational perspective to take given any body of evidence, it is not in general available to us (White 2009: §3). Whether that is viable or not, everyone can agree that the actual provision of a uniquely rational prior would show permissivism to be false, so I will concentrate on constructive proposals in what follows. Proponents of particular constructive projects are known as *objective Bayesians*, and so I will focus on the rivalry between Bayesian permissivists and extant objective Bayesians in what follows. I wish to resist the appropriation of the terminology of 'objective Bayesianism' by proponents of uniqueness however; the permissivist theory of epistemic perspectives defended in Section 3 is not a subjectivist account, but it is compatible with permissivism. I make no secret below of the fact that I have a great deal of sympathy for the project of permissivist objective Bayesianism.

5.2 Permissivism and Priors

One popular early broadly permissivist approach was to try and argue that while permissivism was true of 'informationless' priors, all such priors end up converging to a Unique shared conditional probability when given the same evidence – the priors *wash out*, as it is sometimes put: 'empirical evidence will bring together any two points of view provided they are not dogmatic with respect to each other' (Gaifman and Snir 1982: 498). The mathematical elegance of these convergence-of-opinion theorems is undeniable, but they have strong assumptions and rather weak conclusions (Earman 1992: 141–54). The requirement that the perspectives to be merged not be dogmatic with respect to each other requires that they assign probability zero to the same hypotheses. In the absence of Regularity, and in the presence of very rich spaces of possible hypotheses, very many pairs of acceptable perspectives will therefore not meet the preconditions to be reconcilable with one another. The convergence results also give no indication of the time frame for the priors to

about evidential standards is practically inert, because individual rationality requires deference to peers in a way that secures credal agreement.

wash out, rendering them ineffective as an explanation of current scientific consensus on evidential support.

Recalling that evidential support and confirmation have been defined relative to an epistemic perspective (evidential probability model), the permissivist can argue that rationality and objectivity have been secured. It is an objective fact that $\Pr(H \mid E) > \Pr'(H \mid E)$, for suitable Pr and Pr′; that is the kind of fact that inductive logic yields. The choice as to whether Pr and Pr′ ought to feature in scientific *inference* is a matter that goes beyond inductive logic.

> One might respond at this point by asking, Where do the probability models M come from? and how does one choose an 'appropriate' probability model in a given inductive logical context? These are good questions. However, it is not clear that they must be answered by the inductive logician *qua* logician.... It is not the business of the inductive logician to tell people which probability models they should use (presumably, that is an epistemic or pragmatic question), but once a probability model is specified, the inductive logical relations in that model ... are determined objectively and non-contingently. In the present approach, the duty of the inductive logician is (simply) to explicate the [confirmation]-function—not to decide which probability models should be used in which contexts. (Fitelson 2005: 391–92; cf. Earman 1992: 159)

One might follow the discussion of §4.2 and appeal to context as supplying evidential standards, just as it supplies other parameters to complete overtly unsaturated expressions. It has been argued that natural language quantifiers like 'every' and 'some' must involve reference to domains of quantification, supplied automatically by context when no overt domain is specified (Stanley 2000; Stanley and Szabó 2000). While epistemic contextualism about 'knows' is fiercely contested, that 'is confirmed' or 'is supported' are gradable adjectives is quite plausible. On the present view natural language uses of 'supports' or 'confirms' will pick up some contextually supplied probability model in order to have a semantic content at all; it is unsurprising that in the course of a single conversation the same model will be supplied for all occurrences of 'E confirms H' where no epistemic perspective is explicitly mentioned. The contextualist approach needn't require that any explicit calculation take place to generate an appropriate probability model; it will be one that 'fits' the general background beliefs of speakers and makes the claims they make about confirmation and evidence broadly true. (Again accommodation will play a role here: say '*E* supports *H*' and you thereby make your context one in which the relevant standards make that true, other things being equal.)

Invoking context may seem to have all the advantages of theft over honest toil. So the permissivist may wish to be more specific about why some epistemic perspective might be a candidate for contextual selection. Permissivists have a story to tell about this. It will be one that, like Hume's own account of induction, aims to explain where it cannot justify. There is no pre-given ideal to which we must conform; the explanation of our shared epistemic standards then must appeal to factors which might plausibly produce the phenomenon.

Hume appealed to both 'custom or habit'(1777: ¶5.5) and 'instinct' (1777: ¶9.6), and doubtless both, in updated forms, may play a role in explaining our inductive practice. (Nowadays we might well explain instinct in its turn as the product of natural selection.) The inclination towards having certain priors that produce relatively swift 'learning from experience' (or perhaps, 'jumping to inductive conclusions') is certainly evident in practice. The rational critique of such priors will generally proceed not from selecting some other prior *a priori*, but selecting some rival prior, more cautious or responsible, that resembles the hasty prior in many ways. (Perhaps it will be one that takes the same evidence to be confirmatory of the same hypotheses, but where the degree of confirmation is lower, and hence to approach to inductively-based confidence in a generalisation will be slower.) The point is that the scientific method might involve a refinement of our habits, not the heroic creation of a theory of evidential support out of whole cloth.

Another factor must be sociological – Hume's 'custom'. Scientists are trained, not born. They are enculturated into the scientific mindset, learning through exposure to their mentors and the literature which hypotheses are seen as viable, what sort of evidence is taken to provide a compelling test, etc. If the scientific method can be captured by some constraints on epistemic perspectives, and those constraints are widely endorsed, and there is considerable benefit to being in line with community opinion on confirmation (as there is in actual scientific communities), that is a prudential reason for budding scientists to respect those constraints in the evaluation of evidence.

The permissivist who appeals to sociological or instinctual factors does open themselves up to a charge of arbitrariness (White 2005: 451–52; Feldman 2007: 204–5). Had the background factors been different – had you been differently trained – different epistemic perspective would have been open to you to adopt. While your current standards suggest that E is evidence for H, other standards you could easily have had (had you gone to graduate school elsewhere and had a different mentor) would suggest that E undermines H. Suppose you chose your graduate school for epistemically irrelevant reasons.[2] Can you really think your current attitudes about evidential support are defensible given their fragility?

But the counterfactual about evidential support in no way suggests that you have to dissociate yourself from your current standards once you acknowledge permissivism. What it is to adopt some standards as your own is to regard them as conducing to rational belief. If your response to cases is to be open to thinking those standards might not be reliable, then one hasn't fully adopted those standards. Once one has adopted them, however, one is committed *ex cathedra* to judging that other standards are defective. After all, while H is likely to be true given E, those other standards say it is likely to be false! So those standards will probably get things wrong. One might, as a permissivist, think that one is *lucky* to have been trained in such a way as to have reliable standards, unlike one's peers elsewhere, who are rational but unlucky. But one cannot take their rationality to be a reason to abandon reliable standards, either by suspending judgment on the verdicts of one's own standards, or plumping for rival standards. There is no standpoint-independent 'meta-perspective' that gives one neutral standards for evaluating epistemic standards (Schoenfield 2012: 202; cf. Horowitz 2014: 42–45); there is only where you are.[3]

Whatever the merits of this response to the worry about arbitrariness, it

2 Perhaps, like me, you wanted to be close to New York.

3 Similar things might be said about other standards – perhaps the right thing to say about aesthetic standards is broadly permissivist, but acknowledging that others can be rational in deploying different aesthetic standards doesn't require you to change your evaluation of artwork.

remains unsettling to think that scientific rationality could involve any element of luck or convention. And a compelling answer to any sort of permissivism would be the provision of a rational prior that supported our inductive practice while meeting the desiderata of neutrality and non-arbitrariness implicit in the problem of the priors. In the remainder of this section, I will consider a number of attempts to carry out this task.

5.3 Constructing Priors: the Principle of Indifference

All prominent attempts to construct neutral priors take as their starting point the Leibnizian idea that probability is graded possibility. The uniquely best measure of the degree of possibility – the best probability function – is the one that reflects the natural structure of the space of possibilities. Various proposals have been offered that claim to discern this natural structure. The *classical theory* of probability is a good place to start.

> The theory of chance consists in reducing all the events of the same kind to a certain number of cases equally possible, that is to say, to such as we may be equally undecided about in regard to their existence, and in determining the number of cases favorable to the event whose probability is sought. The ratio of this number to that of all the cases possible is the measure of this probability, which is thus simply a fraction whose numerator is the number of favorable cases and whose denominator is the number of all the cases possible. (Laplace 1825: 6–7)

The classical theory says that 'equal possibilities' should be assigned equal probabilities, and that every probability is reducible to some combination of equal probabilities. The theory was presented as an account of physical probability. It was inadequate to that task, as it could not handle infinite outcome spaces, and excluded the possibility of basic cases with unequal probabilities, such as a biased die.[4] But it is more promising as an account of prior probability. Laplace talks of cases about which 'we *may* be equally undecided about

4 Or, if 1, 2, etc., are not basic cases in the case of a weighted die, then what are the basic cases that allow, e.g., a $1/5$ chance of getting a 6?

in regard to their existence' (my emphasis), and this can be read as suggesting a lack of evidence either way. In that case, Laplace is offering an early version of:

Principle of Indifference (POI) 'if there is no *known* reason for predicating of our subject one rather than another of several alternatives, then relatively to such knowledge the assertions of each of these alternatives have an *equal* probability' (Keynes 1921: 42; cf. White 2009: §§1–2).

The POI is more modest than the classical theory, because it doesn't purport to assign probabilities to all outcomes. It is also a principle that takes an explicitly epistemic attitude, of indifference between possibilities, and yields a determinate probability distribution. The POI, as constructed, is designed to ensure the neutrality of initial probabilities over hypotheses. It works well in many toy examples. What is the rational probability that a goat is behind a given door of the three before you, if two of them have a goat? You have no reason to suppose a goat is behind any particular door; you shouldn't be more confident for no reason, so you should be as neutral as possible, assigning 2/3 probability to each proposition of the form 'a goat is behind door *n*'. This is an implementation of Uniqueness, because POI says the uniquely rational perspective in a situation of equipollence is the indifferent one.

From the perspective of inductive logic, this enforcement of neutrality makes other problems of the priors worse. For the indifferent prior distribution seems to make very poor predictions about inductive support. A binary A/B process of unknown bias will occur 9 times (Weisberg 2011: 507). You have no reason to think it fair; nor to think it biased; nor, if biased, that the bias is in any specific direction. The POI mandates, it seems,[5] a uniform distribution over each hypothesis, i.e., over all 2^9 possible outcome sequences. What is the conditional probability that all outcomes are Bs, given the first 8 are Bs?

$$\Pr(9\text{ Bs} \mid 8\text{ Bs}) = \frac{\Pr(\text{BBBBBBBBB})}{\Pr(\text{BBBBBBBBB} \vee \text{BBBBBBBBA})} = \frac{1/2^9}{2/2^9} = \frac{1}{2}.$$

5 Keynes (1921: ch. 4) suggests that in a case of unknown bias like this, ideal rationality forbids any numerical assignment of probability; that saves the POI at the cost of drastically reducing its concrete role in fixing the priors.

The POI has generated not just uniformity over hypotheses about outcomes prior to experience, but also posterior to experience. Given a process of previously unknown bias, inductive plausibility strongly suggests that 8 (or 88, or 888, ...) consecutive Bs is strong support for a particular hypothesis about bias. We should at least be able to adopt priors that *permit* us to ignore anti-inductive hypotheses, such as regarding BBBBBBBBA as less plausible than BBBBBBBBB. But indifference mandates that we give undue regard to such hypotheses. Neutrality trumps inductive plausibility.

Some might challenge this. In a process of unknown bias, we ought to be indifferent not over individual sequences, but over the *frequencies* those sequences exhibit. Then we should be indifferent over the space of hypotheses '9 As', '8 As, 1 B', ..., '9 Bs'. We will return to the merits of this particular proposal below (§5.4), but there is a worry before we even work out the details: what mandates our representing the problem this way, rather than the first way? This brings out the fact that every application of the POI – just as the classical theory does – a classification of the space of all possibilities into 'basic cases' or 'alternatives' between which we are indifferent. We have to do this; there are so many possible worlds that the only indifference measure over them assigns every possibility no probability at all. So we have to partition the space of possible worlds into basic alternatives in order to get a non-trivial indifference measure. But there are different ways of carving up the very same possibilities (Meacham 2014: 1193–98). Consider this example.

Mystery Cube (van Fraassen 1989: 303) A tool factory produces metal cubes with edge length x, where $1 < x \leq 3$\$. What is the probability that a cube has edge length ≥ 2 cm, given that it was produced by that factory?

The issue is that there are logically equivalent ways of dividing up the same possibilities which seem to give different answers. For we could also represent the possible outputs of the cube factory in terms of their face area, or their volume. Let L be the proposition 'a cube has side length ≥ 1'. The possible cases, and favourable-to-L cases, are detailed in table 5.1.

Table 5.1: Applying the POI to different partitions of Mystery Cube.

Possible cases	Favourable-to-L	Pr(L)
edge length $\in [1, 3]$	length $\in [2, 3]$	$1/2$
face area $\in [1, 9]$	area $\in [4, 9]$	$5/8$
cube volume $\in [1, 27]$	volume $\in [8, 27]$	$19/26$

Apply the POI naively, and we get inconsistent probability assignments. Hesitate to apply it, on the grounds that different partitions give rise to different judgments of epistemic symmetry (White 2009: §3), and we get no probability assignment at all, even though the POI 'is supposed to fill the gap left by missing information' (van Fraassen 1989: 304).

One might have the sense that the POI has been applied incorrectly. In the mystery cube case, we have a problem with multiple representations. Given a representation, e.g., that areas were between 1 and 9 cm^2, POI was applied to the $[1, 9]$ interval to generate the probabilities. But this is manifestly implausible, since it applies indifference to features of the representation, rather than features of the problem represented. A better model would be to identify which representations are merely 'modes of presentation' of the original problem, using those to define a class of transformations that preserve the structure of the original problem. As Rosenkrantz puts it:

> The needed invariances, however, are not obtained by looking at parameter transformations *per se*, but at transformations of the problem itself into equivalent form. Given the statement of the problem, it may for example, be indifferent in what scale units the data are expressed. Such 'indifference between problems' determines what parameter transformations are admissible – not the other way around. (Rosenkrantz 1977: 63; see also Jaynes 1968: 128)

Then POI must be applied in a way that is invariant under those transformations; in practice, to some measure over $(0, 2]$ cm that is equivalent to $(0, 4]$

cm². In the mystery cube case, the allowable transformations are dilations, so the right measure μ on the intervals is $\mu[x, y] = \log y - \log x$. Then we get the 'right' answer (van Fraassen 1989: 310):[6]

$$\Pr(L) = \frac{\mu[2,3]}{\mu[1,3]}\frac{\log 3 - \log 2}{\log 3 - \log 1} = \frac{2\log 3 - 2\log 2}{2\log 3 - 2\log 1} = \frac{\log 9 - \log 4}{\log 9 - \log 1} = \frac{\mu[4,9]}{\mu[1,9]}.$$

This kind of move requires some substantive knowledge about which formal transformations of descriptions of the space of possibilities are those that preserve the 'essential' symmetries of the problem. So this cannot be a purely neutral ignorance prior. Once we have recognised that most circumstances in which we'd wish to apply POI actually involve some background knowledge, the POI turns out to be inapplicable. But there is a generalization of the POI that might apply:

Uniqueness (Maximum Entropy) Given a set C of probability functions meeting certain constraints imposed by the evidence, the uniquely determined evidential probability in light of that evidence is the $\Pr \in C$ such that $H(\Pr) = -\sum_\omega \Pr(\omega)\log\Pr(\omega)$ is maximised, assuming there is exactly one (Jaynes 1957).

The rationale for the Maximum Entropy principle is that entropy is a measure of uninformativeness; so maximum entropy subject to constraints is a way of maximising neutrality given those constraints [Williamson (2011), §8; seidenfeld-1986]. The Maximum Entropy approach does hold out the prospect, unlike the original POI, of both satisfying our desire for neutrality and our desire to have probability functions that are responsive to potential experience (Williamson 2011: §9).

Unfortunately, it would be too hasty to think this gives us a case for Uniqueness. Whenever the uniform distribution is consistent with the background evidence, it always has maximum entropy. But there is no guarantee, if the constraints rule out the uniform distribution, that there is a unique entropy maximising distribution (Shackel and Rowbottom 2020); maximum entropy may turn out to be a moderate permissivist view. This non-uniqueness, as in the original problem cases for the POI, turns out to

6 Even this fails for some cases where there is no neat class of allowable transformations, e.g., those involving both translation and dilation (Milne 1983).

depend on how the problem scenario is represented. (In Williamson's (2011: §10) approach, this is manifest in an explicit language-relativity – see also Weisberg (2011: 508).)

Ultimately, the POI and Maximum Entropy proposals are plausible because they answer, if we are lucky, the twin demands of Uniqueness and neutrality. But once permissivism is brought into view, a view that generates unique probability distributions doesn't look neutral. For example, the original POI mandated a policy of taking any new evidence to be irrelevant to confirmation; while this may be a permissible attitude, it hardly looks mandatory. The most natural maximum entropy distribution that permits responsiveness to experience is one that determines a very specific rule about *how responsive* to be (Weisberg 2011: 508) – yet, intuitively, there is room for variation in appetites for epistemic risk, from Jamesian boldness to Cliffordian timidity. Enforced neutrality between hypotheses leads to overly determinate prescriptions about responsiveness to evidence. (In the Bayesian framework, the fact that unconditional probabilities of hypotheses are expectations of conditional probabilities given possible evidence – i.e., policies for responding to evidence – yokes these two quantities together.) If instead we are not prescriptive about $\Pr(H \mid E)$, remaining neutral to the extent we can over its value, then we won't be as interested in prescriptivism about $\Pr(H)$, and we can secure indifference, where appropriate, by substantive assumptions about the problem scenario at hand. These observations apply also to the remaining attempts to construct explicit unique priors we will consider.

5.4 Constructing Priors: Carnap's Inductive Logic

Treating Carnap at this point is anachronistic; his contributions to inductive logic really kicked off the field, and everyone working on the topics since is indebted to his framing. But he did offer a particular recipe for constructing unique priors, one that really would – if successful – vindicate the idea of an inductive logic. For just as deductive logic gives us relations on sentences in virtue of logical form, so Carnap proposed to give a purely formal account of evidential support:

> While a statement of statistical probability asserts a matter of fact, a statement of inductive probability is of a purely logical nature. If

> hypothesis and evidence are given, the probability can be determined by logical analysis and mathematical calculation. (Carnap 1955: 3)

Given our preceding discussion it seems the prospects for such a proposal are fairly dim, but it is nevertheless worth going through the details, for completeness sake and because it allows us to bolster some of our earlier conclusions against purely formal treatments of epistemic support (§2.3). It also feeds nicely into our subsequent discussion of algorithmic probability (§5.5) and allows us to touch on some issues, like 'gruesome' predicates, that have been implicit so far.

Suppose we have a predicate language, with the connectives of sentential logic, constant terms, and predicates (leaving quantifiers aside). A very simple language might have a single monadic predicate F, and constants denote successive observations in which F might be observed. Evidence consists in a finite binary sequence indicating the presence or absence of F, and our problem is to figure out what evidence sequences provide support for hypotheses about subsequent observations. An 'inductive method' (Carnap 1955: 10) is a procedure for assigning probabilities to hypotheses about the total sequence of outcomes – what Carnap calls a 'state description'. (Every proposition expressible in this language is a Boolean combination of state descriptions.) We've already seen an example of this sort in the previous section, where we noted that a probability assignment that assigns each finite binary A/B sequence of a given length equal probability fails to be inductively rational.

Carnap – and before him, Johnson (1932) – proposed another inductive method than that of the naive POI. A *structure description* is a class of state descriptions which share the same frequencies; i.e., they share that structural aspect which is preserved under permutation of outcome order. Carnap opts for this account of structure as particularly appropriate for statistical inference, because such structures preserve frequencies, which are vital for probabilistic theories. In our previous case, as noted above (§5.3), there are 10 possible structures of the 2^9 possible state descriptions. Carnap's 'method II' says: we ought to be indifferent between structures first, then states (Carnap 1955: 8–14; Zabell 2011: 271–71). Assign equal probability to each structure, then divide that probability equally over each state compatible with a given structure.

That gives us the probability distribution in table 5.2.

Table 5.2: Carnap's 'Method II' for determining inductive probability.

Structure	Structure Probability	Open State Descriptions	State Probability m^*
9 As, 0 Bs	1/10	1	1/10
8 As, 1 B	1/10	9	1/90
7 As, 2 Bs	1/10	36	1/360
6 As, 3 Bs	1/10	84	1/840
5 As, 4 Bs	1/10	126	1/1260
4 As, 5 Bs	1/10	126	1/1260
3 As, 6 Bs	1/10	84	1/840
2 As, 7 Bs	1/10	36	1/360
1 A, 8 Bs	1/10	9	1/90
0 As, 9 Bs	1/10	1	1/10

Carnap's m^* can be extended to a full probability function on the space of possibilities c^*. This enables us to evaluate the prior probabilities of hypotheses, as well as conditional probabilities of hypotheses given evidence. The prior probability assigned to 9 Bs is $1/10$; the prior probability assigned to $BBBBBBBBA$ is $1/90$. So the probability that the last outcome is a B, given that the 8 preceding outcomes have been Bs, is $\frac{1/10}{1/10+1/90} = 9/10$. Half of all states terminate in a B, so the prior probability is $1/2$; so the observation of 8 Bs strongly supports the hypothesis that the last item will be a B, and confirms it over its initial probability. This method does allow for responsiveness to potential evidence. And it does give uniqueness: given a language, purely formal *syntactic* features of state and structure descriptions yield a probability assignment to all hypotheses.

It's bound to be too good to be true. How could syntactic considerations determine a probability assignment over propositions, when the very same proposition can be expressed by sentences with differing syntactic structure? Inconsistency seems unavoidable.

Suppose some speakers have introduced a word, 'grue', the usage of which turns out to best systematised by the following: something is grue iff it is green and examined before December 31, 2030, or blue and not examined before that date (Goodman 1954: ch. 3).[7] At the time of writing, and probably of your reading, everything is green iff it is grue. Suppose we've seen a green/grue regularity in the data so far. If we are to learn from experience in line with Carnap's c^*, we should be confident that green things are grue, and vice versa, going forward. But the first green thing observed on January 1, 2031 will not be grue. So this inductively supported expectation will not be fulfilled.

Carnap's theory cannot accommodate this fact, because the syntactic form of the 'grue' hypotheses and observations is exactly the same as that of the parallel hypotheses and observations including 'green'. 'Grue' was introduced on the basis of a false inductive hypothesis, but having been introduced, it is a fit predicate for use in the construction of state descriptions and structure descriptions, and for the construction of a rational prior. The problem with 'grue' arises once we look at the consequences of applying this syntactic procedure, together with our grasp on the meaning of 'grue'.

To avoid this consequence, we shall have to appeal to some syntactically available feature to exclude 'grue' from our inductive practice. It will be very hard to do so, at least without making unwarranted presumptions about the range of permissible hypotheses (Godfrey-Smith 2003: 578–83). Many things behave differently when observed: people and other social entities, certainly. So there cannot be a general ban of mentioning 'observation' in the hypotheses we consider. Likewise, some data series exhibit discontinuities due to a change in measurement procedure on a certain date; hypotheses that account for the data need to explicitly recognise that date in explaining the slightly different characteristics of the data before and after it. So there cannot be a general ban on mentioning specific dates in hypotheses. 'Grue' is distinctive in including both of these non-forbidden expressions, but we can

7 Imagine perhaps a religious cult, convinced of some puzzling views about the corrupting power of the human gaze, and the idiosyncratic preferences of their deity for blue and uncorrupted things, or green and corrupted things. An anthropologist, sick of writing 'green and corrupted by the human gaze, or blue and not corrupted by the human gaze' might introduce the convenient term 'grue' in describing the cultist's practice; that usage doesn't depend on accepting their dogma, or thinking it any way plausible.

perhaps see a use case for a similar predicate – e.g., when a social science data series changes from using an overt to covert measurement technique, changing the impact of observation on the experimental subjects from a particular date.

The problem with 'grue' isn't intrinsic to the word. The problem is what it *means*. Our background knowledge suggests to us that 'grue' is not a good way of describing reality - it is not a property that captures the structure of how things are. For example, unlike a genuine property, satisfying 'grue' needn't make for genuine resemblance between things which do so (Lewis 1983: 345). But this distinction in 'naturalness' between 'green' and 'grue' isn't – cannot be – present in the syntax. The flipside is that any robust theory of confirmation and evidential support will have to make initial presumptions that treat structurally identical claims differently:

> A favoring relation that fails to treat [structurally indistinguishable] identically *plays favorites* among properties. That is, it responds differently to a hypothesis involving one property than it does to a hypothesis that is identical except that it involves a different property. For instance, suppose we have a piece of evidence that mentions greenness and grueness in exactly the same ways, but that evidence favors a hypothesis involving the property of being green over a hypothesis that involves the property of being grue in structurally identical ways. If the evidential favoring relation behaves in this way, it fails to treat predicate permutations identically. And notice that this property favoritism *precedes the influence of the evidence*. It's not that the difference occurs because the evidence indicates that greenness is a property worthy of special consideration; we stipulated that the evidence says exactly the same things about (or using) greenness that it says about (/using) grueness. If we could behold the [evidential favouring] relation itself before any evidence had been plugged in, we could already see that plugging in evidence and hypotheses involving certain properties would cause it to react differently than plugging in evidence and hypotheses that differed only in the properties that appeared. (Titelbaum 2011: 482–83)

5.5 Constructing Priors: Algorithmic Randomness

Carnap's theory is historically important. But it is also of interest because of the rise of digital computation, where many problems thought to be the exclusive province of human intelligence have shown themselves to be amenable to being handled by systems applying syntactic rules of computation. Carnap does not seem himself to have been particularly interested in formulating an inductive algorithm, but the principal intellectual inheritors of his project in computer science have been. Their story is not as well known among philosophers, so I choose it as my final attempt at explicit construction of a unique prior. It will not be immune to the problems besetting earlier accounts (§§5.3, 5.4). I will simplify some of the mathematical detail in the interest of accessibility (Li and Vitanyi 2008: ch. 2; Eagle 2016b).

Suppose we had some string of sample data about a population and a monadic property F: that the first three items sampled had F, the fourth and fith lacked it, the sixth has it, etc: $F, F, F, \neg F, \neg F, F, \neg F, \neg F, \neg F, \neg F, \neg F, F, \neg F, \neg F, F, \ldots$. Or, as a binary sequence: $1, 1, 1, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, \ldots$. The question we face is one of *sequence prediction* (Solomonoff 1964: 2): what should we say about the next element in this sequence, given the sample so far? Carnap's preferred method took the frequency of Fs into consideration, the balance of F as against ¬F. But it paid no overt attention to the internal structure of the sequence of outcomes.

Laplace observed that, when tossing a coin,

> if heads comes up a hundred times in a row, then this appears to us extraordinary, because the almost infinite number of combinations that can arise in a hundred throws are divided in regular sequences, or those in which we observe a rule that is easy to grasp, and in irregular sequences, that are incomparably more numerous. (Laplace 1825: 16–17)

Laplace notes that an *orderly* sequence is extraordinary if thought to have come about by chance, less extraordinary if it is explicable. As it happens, Carnap's approach respects Laplace's intuition, because it allows orderly data to strongly support orderly hypotheses. The structure descriptions are equiprobable, but there are lots of ways of satisfying the structure 'about half 1s' and only one way of satisfying the structure 'all 1s', so the latter highly orderly sequence gets relatively high probability compared to any one of the many sequences satisfying the former, most of which are random and disorderly.

Carnap has stumbled upon something here: *simplicity*. Orderly sequences obey simple rules; disorderly sequences do not. And since orderly sequences are strongly confirmed by data in agreement with them, we get a preference for simplicity built in to Carnap's inductive methods. But Carnap hasn't latched onto the right mathematical approach. There are still some orderly sequences in the 'about half' structure: $1, 0, 1, 0, 1, 0, \ldots$, for example. This has the frequencies of a disorderly sequence, but isn't disorderly. Carnap's conception is that order is *uniformity*. But in fact, order is exhibited whenever a sequence has a *pattern*. Carnap's framework favours the orderly sequence 'all 1s' over the equally orderly sequence 'alternate 1 and 0'; intuitively, however, that latter pattern is just as indicative of some non-chance theory of the outcomes. This gives the wrong verdict about some straightforward cases. Suppose we'd seen the sequence $0, 1, 0, 1, 0, 1, 0, 1, 0$, and we wonder what comes next. There are slightly more 0s than 1s in the sample data; so c^* slightly favours a prediction of 0 for the next outcome. But the pattern is clear: we ought to predict 1.

A better approach would not overlook regular patterns, including but not limited to uniform outcomes. Solomonoff (1964: 3) made a bold proposal: to get a prior that learns from data, assign prior probabilities to sequences that are inversely proportional to how much internal *complexity* they have. A complex sequence doesn't have readily theorized regularities; a simpler sequence is amenable to a theoretical explanation. Solomonoff argues that the best prediction of the next in a sequence of outcomes is that outcome which would make the resulting sequence simpler.[8]

Solomonoff implements his proposal by linking complexity to *compressibility* (Kolmogorov 1963; Li and Vitanyi 2008: 339–70). Fix on a general purpose computable function f that maps certain binary input sequences to binary output sequences. When $f(\delta) = \sigma$, say that δ is an f-description of σ, or that f decodes δ into its unencoded form σ. A sequence is compressible to

8 This is a mirror image of the 'best systems' analysis of laws of nature, which proposes that laws are those regularities that most simply and powerfully systematise the pattern of events (Lewis 1994: 480).

the extent that the length of its shortest f-description is considerably shorter than it. This can be used to construct a probability function that favours compressible hypotheses, which is to say, it favours hypotheses that posit orderly structure over those that posit randomness. Where '$|x|$' is the length of x, the f-probability of σ is defined (Solomonoff 1997: 2; Rathmanner and Hutter 2011: 1119–21):

$$\Pr_f(\sigma) \stackrel{\text{def}}{=} \sum_{\delta_i \in \{\delta : f(\delta)=\sigma\}} 2^{-|\delta_i|} \approx 2^{-\min\{|\delta| : f(\delta)=\sigma\}}.$$

The f-probability of a sequence is determined by the overall brevity of sequences encoding it, which is dominated by the shortest encoding. The *Kolmogorov complexity* $C_f(\sigma)$ is the length of the shortest input to f encoding σ, so the f-probability of σ is approximately $2^{-C_f(\sigma)}$. Because the probability that a binary sequence of length l is produced by an binary random process is 2^{-l}, we can say that while disorderly sequences, which cannot be compressed, have an algorithmic probability roughly equal to the probability they were produced by chance, orderly sequences have a probability very much greater than the probability they were produced by chance; they are, under this measure over the space of all outcomes, substantially favoured by the prior. But this is still a probability function, summing to 1 over all hypotheses; this shows the orderly sequences must be very scarce.

There is a potentially troubling dependence on f here, but Kolmogorov (and Solomonoff) show that there is a universal or 'asymptotically optimal' function μ such that

$$\forall f \exists k_f \forall \sigma C_\mu(\sigma) \leqslant C_f(\sigma) + k_f.$$

Given that k_f is chosen independently of the sequences, for all sequences beyond a certain length, most decoding functions broadly agree: $C_\mu(\sigma) \approx C_f(\sigma)$.

Algorithmic probability has many desirable qualities. From the perspective of induction, prediction using it can be shown to converge to the 'real' probabilities generating a sequence, in the sense that the distance between the algorithmic posterior probability given the evidence and the real hypotheses converges to zero, so long as the true hypothesis has a non-zero prior (Solomonoff 1997: 11; Ortner and Leitgeb 2011: 736; Rathmanner and Hutter 2011: 1124–25). If there is a rule generating the outcomes, the prior bias of algorithmic probability towards hypotheses that invoke a rule to explain a given sequence of outcomes leads the observation of rule-governed outcomes to quickly favour the hypothesis that computably generates those outcomes.

> Early attempts to justify [algorithmic probability] were based on heuristic arguments involving Occam's razor, as well as many examples in which it gave reasonable answers. At the present time, however, we have a much stronger justification than any heuristic argument. [Algorithmic probability] is the only induction technique known to be complete. By this we mean that if there is any describable regularity in a body of data, [algorithmic probability] will discover it using a relatively small sample of the data. (Solomonoff 1997: 2)

This result might seem striking. However, Sterkenburg argues that assuming the true hypothesis has a non-zero prior is a very strong assumption, for it requires the true hypothesis to mirror the inductive assumptions that go into constructing the prior. The conditions on the convergence theorem assume, in fact, that the true hypothesis is equivalent to a particular strategy of making predictions on the basis of data, and that the inductive presuppositions about which hypotheses to favour are just the ones it also adopts. Then convergence isn't so surprising: if nature produces outcomes by deploying a function it has induced from the data in the way we would, then our 'discovering' that function by induction is hardly surprising. 'We got out what we put in, after all' (Sterkenburg 2016: 476). Moreover, the convergence result might seem in danger of showing too much. For finite data there is always a regularity, even if the mechanism is completely random. The hypotheses that a sequence lacks inductive regularity is antecedently disfavoured; but surely neutrality between hypotheses requires that we are open to learning a sequence is incompressible?

Even if its justification on the basis of convergence results was too good to be true, perhaps we can still adopt algorithmic probability as a prior because we think the assumption that nature will obey a computable rule is a reasonable one. The theory then promises to turn an inductive assumption about computability into a prior we could deploy.

The first piece of bad news about this idea is that algorithmic probability isn't itself computable; we can't determine how complex a description of the total data is, so we cannot determine what the prior probability of the hypothesis which predicts it with certainty is.[9]

The second piece of bad news is that the robustness of the compressibility results, designed to reassure us that dependence on a particular decoding algorithm was inessential, is practically a major difficulty. The constant k_f by which each decoder differs from the universal decoder is arbitrarily large – in effect, it may be thought of as instructions telling the universal machine how to pretend to be f. While these instructions may be small in the limit as sequences to be decoded grow arbitrarily, it may completely swamp sequences on the scale we normally treat, so that choice of decoding function will matter a great deal for particular applications. Sterkenburg (2016: 472–74) argues that the choice here is like the choice of prior for the permissive Bayesian – the existence of the universal machine is like the convergence of opinion theorems (§5.2), and like them gives no guarantee that the convergence is quick enough for practical use.

> while Solomonoff's framework … may offer a theoretical solution to the problem of induction, it cannot be directly applied to practical problems. (Ortner and Leitgeb 2011: 736)

The framework is also subject to more general worries that are at some distance from the technical details. Most obviously, a sequential prediction framework cares only about patterns in the observed data that permit useful

[9] Suppose you could compute $C_\mu(\sigma)$: maybe you enumerate all the sequences of length shorter than σ, feed them sequentially into μ, and see if σ pops out. But sometimes μ won't halt on a given input, so this isn't an effective procedure; *for all you know* you have input a short description of σ and you are just waiting for an unboundedly long decoding process to conclude. There are computable approximations that converge to Kolmogorov complexity; but the error at each stage of convergence is unknown, so they are not approximate in the normal sense of 'close' (Rissanen 1997; Solomonoff 1997: 2).

compression. The sequence of successive green emeralds is no more or less compressible than the sequence of successive grue emeralds, because only the form of the data matters. So if we want to avoid inconsistent predictions, we will have to impose some 'green'-favouring constraint before applying Solomonoff's recipe. Solomonoff induction, no less than Carnpian inductive methods, satisfy the pre-conditions for generalised language-relativity results (Titelbaum 2011: 482–84). This should temper some of the bolder claims made about the approach, e.g., 'through Solomonoff, … the problem of formalizing optimal inductive inference is solved' (Rathmanner and Hutter 2011: 1078).

Solomonoff himself seems to have come to appreciate this sort of observation, about *grue* and about the prior choice of hypotheses to be considered, and connected it to the practical differences between different universal hypotheses:

> choosing a reference machine we are given the opportunity to insert into the *a priori* probability distribution any information about the data that we know before we see the data. (Solomonoff 1997: 4)

There is another difficulty: how do we turn events distributed across space and time into a sequence of data that can be input to Solomonoff induction?

> Suppose that I am tossing a coin on a train that is moving back and forth on tracks that point in a generally easterly direction. … Moving from left to right (west to east), we see the pattern: HTHTHTHTH…. Moving upwards (earlier to later), we see the pattern: HHTHHTHHT…. Imagine, as we can, that these patterns persist forever. What is the limiting relative frequency of Heads? Taking the results in their temporal order, the answer is 2/3…. But taking them in their west-east spatial order, the answer is 1/2. Now, why should one answer have priority over the other? In other words, we have more than one limiting relative frequency, depending on which spatio-temporal dimension we privilege. (Hájek 2009: 218–19)

If the events are spread out sufficiently over spacetime, there may be no determinate frame-invariant fact about temporal order of outcomes (Maudlin

2012: ch. 5), even if induction is understood as essentially temporal in a way we have no presumed here. Carnap's maligned C^* does a little better here, at least for for finite sequences, because the statistical properties on which he relies are invariant under data permutations. Assumptions about the structure of data must clearly come before one can formulate a space of hypotheses about sequences of data, and is substantively dependent in this example on non-formal views about space and time.

The upshot, as I see it, of §§5.3, 5.4, 5.5 is that substantive philosophical assumptions about prior probabilities are unavoidable. The hope that there is only one rational way to think about evidential support, even though it cannot be articulated in any generally compelling way, starts to look more like an article of faith than a reasonable guess. Technical advances may help us formulate prior assumptions, and distributions derived from maximum entropy or Kolmogorov complexity may well have appealing properties as explications of our preference for neutrality or simplicity. But they cannot function without prior assumptions, and do not let us avoid making them. When the mathematical complexity goes up, and the results start looking more like magic, assumptions are no less present than in the case of permissivism, just slightly better disguised. This lesson generalizes well beyond this area of philosophy.

References

Alchourrón, Carlos, Peter Gärdenfors, and David Makinson (1985) 'On the Logic of Theory Change: Partial Meet Contraction and Revision Functions', *Journal of Symbolic Logic* **50**: 510–30.

Armendt, Brad (1993) 'Dutch Books, Additivity and Utility Theory', *Philosophical Topics* **21**: 1–20.

Arntzenius, F, A Elga, and J Hawthorne (2004) 'Bayesianism, Infinite Decisions, and Binding' **113**: 251–83. doi:10.1093/mind/113.450.251.

Arpaly, Nomy (2011) 'Open-Mindedness as a Moral Virtue', *American Philosophical Quarterly* **48**: 75–85.

Babbage, Charles (1827) 'Notice Respecting Some Errors Common to Many Tables of Logarithms', *Memoirs of the Royal Astronomical Society* **3**: 65–67.

Bacon, Francis (1620/2000) *The New Organon*, Cambridge Texts in the History of Philosophy, Lisa Jardine and Michael Silverthorne, eds. Cambridge University Press.

Baker, Victor R (2009) 'The Channeled Scabland: A Retrospective', *Annual Review of Earth and Planetary Sciences* **37**: 393–411. doi:10.1146/annurev.earth.061008.134726.

Bar-Hillel, Maya and Avishai Margalit (1972) 'Newcomb's Paradox Revisited', *The British Journal for the Philosophy of Science* **23**: 295–304. doi:10.1093/bjps/23.4.295.

Barnes, Luke A (2020) 'A Reasonable Little Question: A Formulation of the Fine-Tuning Argument', *Ergo* **6**: 1220–57. doi:10.3998/ergo.12405314.0006.042.

Bartha, Paul (2004) 'Countable Additivity and the de Finetti Lottery', *British Journal for the Philosophy of Science* **55**: 301–21. doi:10.1093/bjps/55.2.301.

Bovens, Luc and Stephan Hartmann (2003) *Bayesian Epistemology*. Oxford University Press.

Bradley, Darren (2015) *A Critical Introduction to Formal Epistemology*. Bloomsbury Academic.

Bradley, Darren and Hannes Leitgeb (2006) 'When Betting Odds and Credences Come Apart: More Worries for Dutch Book Arguments', *Analysis* **66**: 119–27. doi:10.1111/j.1467-8284.2006.00598.x.

Brössel, Peter and Franz Huber (2015) 'Bayesian Confirmation: A Means with No End', *The British Journal for the Philosophy of Science* **66**: 737–49. doi:10.1093/bjps/axu004.

Buchak, Lara (2013) *Risk and Rationality*. Oxford University Press.

Camp, Elisabeth (2019) 'Perspectives and Frames in Pursuit of Ultimate Understanding', in Stephen R Grimm, ed., *Varieties of Understanding: New Perspectives from Philosophy, Psychology, and Theology*: 17–45. Oxford University Press.

Carnap, Rudolf (1947) 'On the Application of Inductive Logic', *Philosophy and Phenomenological Research* **8**: 133. doi:10.2307/2102920.

Carnap, Rudolf (1955) 'Statistical and Inductive Probability', in *I. Statistical and Inductive Probability, II. Inductive Logic and Science*: 1–16. Galois Institute of Mathematics and Art.

Carnap, Rudolf (1962) *Logical Foundations of Probability*, 2nd edition. University of Chicago Press.

Carr, Jennifer Rose (2021) 'Why Ideal Epistemology?', *Mind* **131**: 1131–62. doi:10.1093/mind/fzab023.

Climenhaga, Nevin (2023) 'Epistemic Probabilities Are Degrees of Support, Not Degrees of (Rational) Belief', *Philosophical and Phenomenological Research*. doi:10.1111/phpr.12947.

Crupi, Vincenzo (2021) 'Confirmation', in Edward N Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2021/entries/confirmation/.

de Finetti, Bruno (1937/1964) 'Foresight: Its Logical Laws, Its Subjective Sources', in Howard E Smokler and Henry E Kyburg Jr, eds., Henry E Kyburg Jr, trans., *Studies in Subjective Probability*: 93–158. Wiley.

de Finetti, Bruno (1974) *Theory of Probability*, vol. 1, A Machí and A Smith, eds. Wiley.

Dorling, Jon (1979) 'Bayesian Personalism, the Methodology of Scientific Research Programmes, and Duhem's Problem', *Studies in History and Philo-*

sophy of Science **10**: 177–87. doi:10.1016/0039-3681(79)90006-2.

Drexel, J F, W V Preiss, and A J Parker, eds. (1993/2012) *The Geology of South Australia*, vol. 1, Bulletin of the Geological Survey of South Australia, no. 54. Government of South Australia. https://sarigbasis.pir.sa.gov.au/WebtopEw/ws/samref/sarig1/image/DDD/BULL054(V1).pdf.

Eagle, Antony, ed. (2011) *Philosophy of Probability: Contemporary Readings*. Routledge.

Eagle, Antony (2016b) 'Probability and Randomness', in Alan Hájek and Christopher R Hitchcock, eds., *Oxford Handbook of Probability and Philosophy*: 440–59. Oxford University Press.

Eagle, Antony (2016a) 'Probability', in Paul Humphreys, ed., *The Oxford Handbook of Philosophy of Science*: 417–39. Oxford University Press.

Earman, John (1992) *Bayes or Bust?*, A Bradford Book. MIT Press.

Eells, E and B Fitelson (2002) 'Symmetries and Asymmetries in Evidential Support', *Philosophical Studies* **107**. doi:10.1023/a:1014712013453.

Eells, Ellery and Branden Fitelson (2000) 'Measuring Confirmation and Evidence', *Journal of Philosophy* **97**: 663–72.

Elga, Adam (2007) 'Reflection and Disagreement', *Noûs* **41**: 478–502. doi:10.1111/j.1468-0068.2007.00656.x.

Feldman, Richard (2007) 'Reasonable Religious Disagreements', in Louise Antony, ed., *Philosophers Without Gods*: 194–214. Oxford University Press.

Fisher, R A (1935) 'The Logic of Inductive Inference', *Journal of the Royal Statistical Society* **98**: 39–54.

Fitelson, Branden (1999) 'The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity', *Philosophy of Science* **66**: S362–78. doi:10.1086/392738.

Fitelson, Branden (2005) 'Inductive Logic', in Jessica Pfeifer and Sahotra Sarkar, eds., *Philosophy of Science: An Encyclopedia*: 384–93. Routledge.

Fitelson, Branden (2007) 'Likelihoodism, Bayesianism, and Relational Confirmation', *Synthese* **156**: 473–89. doi:10.1007/s11229-006-9134-9.

Fitelson, Branden and Alan Hájek (2014) 'Declarations of Independence', *Synthese* **194**: 3979–95. doi:10.1007/s11229-014-0559-2.

Fitelson, Branden and James Hawthorne (2010) 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens', in *The Place of Probability in Science*: 247–75. Springer Netherlands.

Gaifman, Haim (1988) 'A Theory of Higher Order Probability', in Brian Skyrms and William L Harper, eds., *Causation, Chance, and Credence*: 191–219. Kluwer.

Gaifman, Haim and Marc Snir (1982) 'Probabilities over Rich Languages, Testing and Randomness', *The Journal of Symbolic Logic* **47**: 495–548.

Gandenberger, Greg (2016) 'Why I Am Not a Likelihoodist', *Philosophers' Imprint* **16**: 1–22.

Garber, Daniel (1983) 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory', in John Earman, ed., *Testing Scientific Theories*. University of Minnesota Press.

Gärdenfors, Peter (1988) *Knowledge in Flux*. MIT Press.

Glymour, Clark (1981) *Theory and Evidence*. University of Chicago Press.

Glynn, Luke (2010) 'Deterministic Chance', *The British Journal for the Philosophy of Science* **61**: 51–80. doi:10.1093/bjps/axp020.

Godfrey-Smith, Peter (2003) 'Goodman's Problem and Scientific Methodology', *Journal of Philosophy* **100**: 573–90.

Good, I J (1961) 'The Paradox of Confirmation (III)', *British Journal for the Philosophy of Science* **11**: 63–64.

Good, I J (1967) 'The White Shoe Is a Red Herring', *British Journal for the Philosophy of Science* **17**: 322–22. doi:10.1093/bjps/17.4.322.

Goodman, Nelson (1954) *Fact, Fiction and Forecast*. Harvard University Press.

Greaves, Hilary (2013) 'Epistemic Decision Theory', *Mind* **122**: 915–52. doi:10.1093/mind/fzt090.

Greco, Daniel and Brian Hedden (2016) 'Uniqueness and Metaepistemology', *The Journal of Philosophy* **113**: 365–95. doi:10.5840/jphil2016113825.

Grice, Paul (1989) 'Logic and Conversation', in *Studies in the Way of Words*: 22–40. Harvard University Press.

Hacking, Ian (1965) *The Logic of Statistical Inference*. Cambridge University Press.

Hacking, Ian (2001) *Introduction to Probability and Inductive Logic*. Cambridge University Press.

Hájek, Alan (1997) '"Mises Redux"-Redux: Fifteen Arguments Against Finite Frequentism', *Erkenntnis* **45**: 209–27.

Hájek, Alan (2003) 'What Conditional Probability Could Not Be', *Synthese* **137**: 273–323. doi:10.1023/b:synt.0000004904.91112.16.

Hájek, Alan (2005) 'Scotching Dutch Books?', *Philosophical Perspectives* **19**: 139–51. doi:10.1111/j.1520-8583.2005.00057.x.

Hájek, Alan (2008) 'Dutch Book Arguments', in Paul Anand, Prasanta Pattanaik and Clemens Puppe, eds., *Oxford Handbook of Rational and Social Choice*: 173–95. Oxford University Press.

Hájek, Alan (2009) 'Fifteen Arguments Against Hypothetical Frequentism', *Erkenntnis* **70**: 211–35. doi:10.1007/s10670-009-9154-1.

Hájek, Alan (2011) 'Triviality Pursuit', *Topoi* **30**: 3–15. doi:10.1007/s11245-010-9083-2.

Hájek, Alan and Christopher Hitchcock (2017) 'Probability for Everyone—Even Philosophers', in Alan Hájek and Christopher Hitchcock, eds., *Oxford Handbook of Probability and Philosophy*: 5–30. Oxford University Press.

Hall, Ned (2004) 'Two Mistakes about Credence and Chance', *Australasian Journal of Philosophy* **82**: 94–112. doi:10.1080/713659806.

Harman, Gilbert (1965) 'The Inference to the Best Explanation', *The Philosophical Review* **74**: 88–95. doi:10.2307/2183532.

Harman, Gilbert (1986) *Change in View*, A Bradford Book. MIT Press.

Hawthorne, James (2005) '*Degree-of-Belief* and *Degree-of-Support*: Why Bayesians Need Both Notions', *Mind* **114**: 277–320. doi:10.1093/mind/fzi277.

Hawthorne, John (2004) *Knowledge and Lotteries*. Oxford University Press.

Hedden, Brian (2015) *Reasons Without Persons*. Oxford University Press.

Hempel, Carl G (1945a) 'Studies in the Logic of Confirmation I', *Mind* **54**: 1–26. doi:10.1093/mind/liv.213.1.

Hempel, Carl G (1945b) 'Studies in the Logic of Confirmation II', *Mind* **54**: 97–121. doi:10.1093/mind/liv.214.97.

Henderson, Leah (2014) 'Bayesianism and Inference to the Best Explanation', *The British Journal for the Philosophy of Science* **65**: 687–715. doi:10.1093/bjps/axt020.

Hendrick, R Edward and Mark A Helvie (2011) 'United States Preventive Services Task Force Screening Mammography Recommendations: Science Ignored', *American Journal of Roentgenology* **196**: W112–16. doi:10.2214/AJR.10.5609.

Horowitz, Sophie (2014) 'Immoderately Rational', *Philosophical Studies* **167**: 41–56. doi:10.1007/s11098-013-0231-6.

Horwich, Paul (1982) *Probability and Evidence*. Cambridge University Press.

Hosiasson-Lindenbaum, Janina (1940) 'On Confirmation', *Journal of Symbolic Logic* **5**: 133–48.

Howson, Colin (1991) 'The "Old Evidence" Problem', *British Journal for the Philosophy of Science* **42**: 547–55.

Howson, Colin (2000) *Hume's Problem: Induction and the Justification of Belief*. Oxford University Press.

Howson, Colin and Peter Urbach (1993) *Scientific Reasoning: The Bayesian Approach*, 2nd edition. Open Court.

Huber, Franz (2008) 'Hempel's Logic of Confirmation', *Philosophical Studies* **139**: 181–89. doi:10.1007/s11098-007-9111-2.

Hume, David (1777/2022) *An Enquiry Concerning Human Understanding*, Amyas Merivale and Peter Millican, eds. https://davidhume.org/texts/e/.

Ismael, Jenann (2008-06) 'Raid! Dissolving the Big, Bad Bug', *Noûs* **42**: 292–307. doi:10.1111/j.1468-0068.2008.00681.x.

Jackson, Elizabeth (2019) 'How Belief-Credence Dualism Explains Away Pragmatic Encroachment', *Philosophical Quarterly* **69**: 511–33. doi:10.1093/pq/pqz006.

Jaynes, Edwin T (1957) 'Information Theory and Statistical Mechanics', *Physical Review* **106**: 620–30. doi:10.1103/PhysRev.106.320.

Jaynes, Edwin T (1968) 'Prior Probabilities', *IEEE Transactions on Systems Science and Cybernetics* **4**: 227–41. doi:10.1109/tssc.1968.300117.

Jeffrey, Richard C (1983) 'Bayesianism with a Human Face', in John Earman, ed., *Testing Scientific Theories*, vol. **X**: 133–56. University of Minnesota Press.

Jeffrey, Richard C (1992) 'Probability and the Art of Judgment', in *Probability and the Art of Judgment*: 44–76. Cambridge University Press.

Jeffrey, Richard C (2008) *Subjective Probability: The Real Thing*. Cambridge University Press.

Jevons, W S (1874) *The Principles of Science*, vol. 1. Macmillan; Co.

Johnson, W E (1932) 'Probability: The Deductive and Inductive Problems', *Mind* **41**: 409–23. doi:10.1093/mind/xli.164.409.

Joyce, James M (2007-08) 'Epistemic Deference: The Case of Chance', *Proceedings of the Aristotelian Society (Hardback)* **107**: 187–206. doi:10.1111/j.1467-9264.2007.00218.x.

Joyce, James M (1998) 'A Nonpragmatic Vindication of Probabilism', *Philosophy of Science* **65**: 575–603. doi:10.1086/392661.

Kadane, Joseph B and Anthony O'Hagan (1995) 'Using Finitely Additive Probability: Uniform Distributions on the Natural Numbers', *Journal of the American Statistical Association* **90**: 626–31. doi:10.1080/01621459.1995.10476556.

Kahneman, Daniel and Amos Tversky (1979) 'Prospect Theory: An Analysis of Decision Under Risk', *Econometrica* **47**: 263. doi:10.2307/1914185.

Kelly, Kevin T (1996) *The Logic of Reliable Inquiry*. Oxford University Press.

Kelly, Kevin T and Clark Glymour (2004) 'Why Probability Does Not Captures the Logic of Scientific Confirmation', in Christopher Hitchcock, ed., *Contemporary Debates in Philosophy of Science*: 94–114. Blackwell.

Kemeny, John G (1955) 'Fair Bets and Inductive Probabilities', *Journal of Symbolic Logic* **20**: 263–73.

Kemeny, John G and Paul Oppenheim (1952) 'Degree of Factual Support', *Philosophy of Science* **19**: 307–24.

Keynes, John Maynard (1921) *A Treatise on Probability*. Macmillan.

Kim, Brian (2017) 'Pragmatic Encroachment in Epistemology', *Philosophy Compass* **12**: e12415 14. doi:10.1111/phc3.12415.

King, Jeffrey C, Scott Soames, and Jeff Speaks (2014) *New Thinking about Propositions*. Oxford University Press.

Kolmogorov, A N (1933/1956) *Foundations of the Theory of Probability*, 2nd edition, Nathan Morrison, ed., trans. Chelsea Publishing Company.

Kolmogorov, A N (1963) 'On Tables of Random Numbers', *Sankhyā* **25**: 369–76. doi:10.1016/S0304-3975(98)00075-9.

Kopec, Matthew and Michael G Titelbaum (2016) 'The Uniqueness Thesis', *Philosophy Compass* **11**: 189–200. doi:10.1111/phc3.12318.

Kyburg, Jr, Henry E (1978) 'Subjective Probability: Criticisms, Reflections, and Problems', *Journal of Philosophical Logic* **7**: 157–80. doi:10.1007/BF00245926.

Kyburg, Jr, Henry E (1983) *Epistemology and Inference*. Univerasity of Minnesota Press.

Laplace, Pierre-Simon (1825/1995) *Philosophical Essay on Probabilities*, 5th edition, Andrew I Dale, trans. Springer New York.

Lehman, Constance D, Robert F Arao, Brian L Sprague, Janie M Lee, Diana S M Buist, *et al.* (2017) 'National Performance Benchmarks for Modern Screening Digital Mammography: Update from the Breast Cancer Surveillance Consortium', *Radiology* **283**: 49–58. doi:10.1148/radiol.2016161174.

Lehman, R Sherman (1955) 'On Confirmation and Rational Betting', *Journal of Symbolic Logic* **20**: 251–62.

Leitgeb, Hannes and Richard Pettigrew (2010b) 'An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy', *Philosophy of Science* **77**: 236–72. doi:10.1086/651318.

Leitgeb, Hannes and Richard Pettigrew (2010a) 'An Objective Justification of Bayesianism I: Measuring Inaccuracy', *Philosophy of Science* **77**: 201–35. doi:10.1086/651317.

Lewis, David (1976) 'Probabilities of Conditionals and Conditional Probabilities', *Philosophical Review* **85**: 297–315.

Lewis, David (1979) 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* **8**: 339–59. doi:10.1007/bf00258436.

Lewis, David (1983) 'New Work for a Theory of Universals', *Australasian Journal of Philosophy* **61**: 343–77. doi:10.1080/00048408312341131.

Lewis, David (1986b) 'Causal Explanation', in *Philosophical Papers*, vol. **2**: 214–40. Oxford University Press.

Lewis, David (1986a) 'A Subjectivist's Guide to Objective Chance', in *Philosophical Papers*, vol. **2**: 83–132. Oxford University Press.

Lewis, David (1994) 'Humean Supervenience Debugged', in *Papers in Metaphysics and Epistemology*: 224–46. Cambridge University Press.

Lewis, David (1996) 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**: 549–67. doi:10.1080/00048409612347521.

Lewis, David (1999) 'Why Conditionalize?', in *Papers in Metaphysics and Epistemology*: 403–7. Cambridge University Press.

Li, Ming and Paul M B Vitanyi (2008) *An Introduction to Kolmogorov Complexity and Its Applications*. Springer Verlag.

Lipton, Peter (2004) *Inference to the Best Explanation*, 2nd edition. Routledge.

Lyell, Charles (1830) *Principles of Geology*, vol. 1. James Murray.

Maher, Patrick (1996) 'Subjective and Objective Confirmation', *Philosophy of Science* **63**: 149–74. doi:10.1086/289906.

Maher, Patrick (1997) 'Depragmatized Dutch Book Arguments', *Philosophy of Science* **64**: 291–305. doi:10.1086/392552.

Maher, Patrick (2002) 'Joyce's Argument for Probabilism', *Philosophy of Science* **69**: 73–81. doi:10.1086/338941.

Maudlin, Tim (2012) *Philosophy of Physics: Space and Time*. Princeton University Press.

McGee, Vann (1999) 'An Airtight Dutch Book', *Analysis* **59**: 257–65. doi:10.1111/1467-8284.00178.

Meacham, Christopher J G (2014) 'Impermissive Bayesianism', *Erkenntnis* **79**: 1185–1217. doi:10.1007/s10670-013-9547-z.

Meacham, Christopher J G (2016) 'Ur-Priors, Conditionalization, and Ur-Prior Conditionalization', *Ergo* **3**. doi:10.3998/ergo.12405314.0003.017.

Milne, Peter (1983) 'A Note on Scale Invariance', *British Journal for the Philosophy of Science* **34**: 49–55. doi:10.1093/bjps/34.1.49.

Milne, Peter (1996) 'log$[P(h|eb)/P(h|b)]$ Is the One True Measure of Confirmation', *Philosophy of Science* **63**: 21–26. doi:10.1086/289891.

Moretti, Luca (2003) 'Why the Converse Consequence Condition Cannot Be Accepted', *Analysis* **63**: 297–300. doi:10.1111/1467-8284.00440.

Morvan, Pierre Le (1999) 'The Converse Consequence Condition and Hempelian Qualitative Confirmation', *Philosophy of Science* **66**: 448–54. doi:10.1086/392697.

Norton, John D (2003) 'A Material Theory of Induction', *Philosophy of Science* **70**: 647–70. doi:10.1086/378858.

Norton, John D (2011) 'Challenges to Bayesian Confirmation Theory', in Prasanta S Bandyopadhyay and Malcolm R Forster, eds., *Handbook of the Philosophy of Science*, vol. **7**: 391–439. Elsevier.

Ortner, Ronald and Hannes Leitgeb (2011) 'Mechanizing Induction', in Dov M Gabbay, Stephan Hartmann and John Woods, eds., *Inductive Logic*, vol. **10**: 719–72. North-Holland.

Paul, L A (2014) *Transformative Experience*. Oxford University Press.

Pettigrew, Richard (2016) *Accuracy and the Laws of Credence*. Oxford University Press.

Pettigrew, Richard (2020a) *Choosing for Changing Selves*. Oxford University Press.

Pettigrew, Richard (2020b) *Dutch Book Arguments*. Cambridge University Press.

Popper, Karl (1959) *The Logic of Scientific Discovery*. Hutchinson.

Pritchard, Duncan, John Turri, and J Adam Carter (2022) 'The Value of Knowledge', in Edward N Zalta and Uri Nodelman, eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. https:

//plato.stanford.edu/archives/fall2022/entries/knowledge-value/.

Pryor, James (2000) 'The Skeptic and the Dogmatist', *Noûs* **34**: 517–49. doi:10.1111/0029-4624.00277.

Quine, Willard van Orman (1951) 'Two Dogmas of Empiricism', *The Philosophical Review* **60**: 20–43.

Ramsey, F P (1926/1990) 'Truth and Probability', in D H Mellor, ed., *Philosophical Papers*: 52–94. Cambridge University Press.

Rathmanner, Samuel and Marcus Hutter (2011) 'A Philosophical Treatise of Universal Induction', *Entropy* **13**: 1076–136. doi:10.3390/e13061076.

Rissanen, Jorma (1997) 'Stochastic Complexity in Learning', *Journal of Computer and System Sciences* **55**: 89–95. doi:10.1006/jcss.1997.1501.

Rosen, Deborah A (1978) 'In Defense of a Probabilistic Theory of Causality', *Philosophy of Science* **45**: 604–13.

Rosen, Gideon (2001) 'Nominalism, Naturalism, Epistemic Relativism', *Philosophical Perspectives* **15**: 60–91. doi:10.1111/0029-4624.35.s15.4.

Rosenkrantz, R D (1977) *Inference, Method and Decision: Towards a Bayesian Philosophy of Science*. D. Reidel.

Royall, R M (1997) *Statistical Evidence: A Likelihood Paradigm*. Chapman; Hall.

Schick, Frederic (1986) 'Dutch Bookies and Money Pumps', *Journal of Philosophy* **83**: 112–19.

Schlesinger, George N (1995) 'Measuring Degrees of Confirmation', *Analysis* **55**: 208–12.

Schoenfield, Miriam (2012) 'Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief', *Noûs* **48**: 193–218. doi:10.1111/nous.12006.

Schupbach, Jonah N (2022) *Bayesianism and Scientific Reasoning*. Cambridge University Press.

Shackel, Nicholas and Darrell P Rowbottom (2020) 'Bertrand's Paradox and the Maximum Entropy Principle', *Philosophy and Phenomenological Research* **101**: 505–23. doi:10.1111/phpr.12596.

Skyrms, Brian (1984) *Pragmatics and Empiricism*. Yale University Press.

Skyrms, Brian (2000) *Choice and Chance*. Wadsworth.

Sober, Elliott (1994) 'Contrastive Empiricism', in *From a Biological Point of View: Essays in Evolutionary Philosophy*: 114–35. Cambridge University

Press.

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

Solomonoff, Ray (1964) 'A Formal Theory of Inductive Inference, Part I', *Information and Control* **7**: 1–22. doi:10.1016/S0019-9958(64)90223-2.

Solomonoff, Ray (1997) *Does Algorithmic Probability Solve the Problem of Induction?* http://world.std.com/~rjs/isis96.pdf.

Staffel, Julia (2020) *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press.

Stalnaker, Robert C (1973) 'Presuppositions', *Journal of Philosophical Logic* **2**: 447–57. doi:10.1007/bf00262951.

Stalnaker, Robert C (1984) *Inquiry*, A Bradford Book. MIT Press.

Stalnaker, Robert C (1998/1999) 'On the Representation of Context', in *Context and Content*: 96–113. Oxford University Press.

Stanley, Jason (2000) 'Context and Logical Form', *Linguistics and Philosophy* **23**: 391–434. doi:10.1023/a:1005599312747.

Stanley, Jason and Zoltán Gendler Szabó (2000) 'On Quantifier Domain Restriction', *Mind and Language* **15**: 219–61. doi:10.1111/1468-0017.00130.

Steel, Daniel (1996) 'Bayesianism and the Value of Diverse Evidence', *Philosophy of Science* **63**: 666–74.

Sterkenburg, Tom F (2016) 'Solomonoff Prediction and Occam's Razor', *Philosophy of Science* **83**: 459–79. doi:10.1086/687257.

Strevens, Michael (2001) 'The Bayesian Treatment of Auxiliary Hypotheses', *British Journal for the Philosophy of Science* **52**: 515–37. doi:10.1093/bjps/52.3.515.

Strevens, Michael (2004) 'Bayesian Confirmation Theory: Inductive Logic, or Mere Inductive Framework?', *Synthese* **141**: 365–79. doi:10.1023/b:synt.0000044991.73791.f7.

Strevens, Michael (2006) 'Notes on Bayesian Confirmation Theory'. http://www.nyu.edu/gsas/dept/philo/user/strevens/Classes/Conf06/BCT.pdf.

Swinburne, Richard (1971) 'The Paradoxes of Confirmation: A Survey', *American Philosophical Quarterly* **8**: 318–30.

Talbott, W J (1991) 'Two Principles of Bayesian Epistemology', *Philosophical Studies* **62**: 135–50. doi:10.1007/bf00419049.

Tarski, Alfred (1983) 'The Concept of Truth in Formalized Languages', in *Logic, Semantics, Metamathematics*: 152–278. Hackett.

Titelbaum, Michael G (2011) 'Not Enough There There Evidence, Reasons, and Language Independence', *Philosophical Perspectives* **24**: 477–528. doi:10.1111/j.1520-8583.2010.00201.x.

Titelbaum, Michael G (2022b) 'Fundamentals of Bayesian Epistemology', in, vol. **2: Arguments, Challenges, Alternatives**. Oxford University Press.

Titelbaum, Michael G (2022a) 'Fundamentals of Bayesian Epistemology', in, vol. **1: Introducing Credences**. Oxford University Press.

van Fraassen, Bas C (1980) *The Scientific Image*. Clarendon Press.

van Fraassen, Bas C (1989) *Laws and Symmetry*. Oxford University Press.

Weisberg, Jonathan (2009) 'Commutativity or Holism? A Dilemma for Conditionalizers', *The British Journal for the Philosophy of Science* **60**: 793–812. doi:10.1093/bjps/axp007.

Weisberg, Jonathan (2011) 'Varieties of Bayesianism', in Dov Gabbay, Stephan Hartmann and John Woods, eds., *Handbook of the History of Logic*, vol. **10**: 477–551. Elsevier.

White, Roger (2005) 'Epistemic Permissiveness', *Philosophical Perspectives* **19**: 445–59.

White, Roger (2009) 'Evidential Symmetry and Mushy Credence', in Tamar Szabó Gendler and John Hawthorne, eds., *Oxford Studies in Epistemology*, vol. **3**: 161–86.

Williams, J Robert G (2016) 'Probability and Nonclassical Logic', in Christopher R Hitchcock and Alan Hájek, eds., *The Oxford Handbook of Probability and Philosophy*: 248–76. Oxford University Press.

Williamson, Jon (1999) 'Countable Additivity and Subjective Probability', *The British Journal for the Philosophy of Science* **50**: 401–16. doi:10.1093/bjps/50.3.401.

Williamson, Jon (2011) 'An Objective Bayesian Account of Confirmation', in Dennis Dieks, Wenceslao J Gonzalez, Stephan Hartmann, Thomas Uebel and Marcel Weber, eds., *Explanation, Prediction, and Confirmation*, vol. **2**: 53–81. Springer.

Williamson, Timothy (2000) *Knowledge and Its Limits*. Oxford University Press.

Williamson, Timothy (2007) 'How Probable Is an Infinite Sequence of Heads?', *Analysis* **67.3**: 173–80. doi:10.1111/j.1467-8284.2007.00671.x.

Zabell, S L (2011) 'Carnap and the Logic of Inductive Inference', *Handbook of* *the History of Logic* **10**: 265–309. doi:10.1016/b978-0-444-52936-7.50008-2.

Zalabardo, J (2009) 'An Argument for the Likelihood-Ratio Measure of Confirmation', *Analysis* **69**: 630–35. doi:10.1093/analys/anp103.